Integrated Machine Learning and Optimization Frameworks with Applications in Operations Management

by

Amirhossein Meisami

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial and Operations Engineering)
in the University of Michigan
2018

Doctoral Committee:

Associate Professor Henry Lam, Co-Chair
Professor Mark P. Van Oyen, Co-Chair
Assistant Professor Cong Shi
Associate Professor Ambuj Tewari

Amirhossein Meisami
meisami@umich.edu
ORCID iD: 0000-0003-3493-4410

To Ava

Acknowledgement

I would like to sincerely thank my advisors, Mark and Henry, for their incredible guidance and inspiration. In particular, I do appreciate the absolute freedom I was given to explore my interests while enjoying tremendous intellectual support.

During the past four years, I have had the privilege to work with several collaborators from St. Joseph Mercy Hospital, Memorial Sloan Kettering Cancer Center and Adobe, Inc. I would like to specifically acknowledge the contribution and support of Mark Cowen, Chris Stromblad, Nick Kastango, Chen Dong and Abhishek Pani.

Further, I like to thank all of my previous instructors from the K. N. Toosi University of Technology, Texas A&M University, and the University of Michigan for their endless efforts inside and outside the classroom. I am also very much thankful to Dr. Kiavash Kianfar for his support and guidance during my two years at Texas A&M University.

I would also like to thank my lovely family, my parents and my sisters, Parisa and Parastoo, for always being there for me. I would not be able to complete this journey without their unconditional love, unlimited support and endless encouragement.

Finally, I dedicate this dissertation to Ava Athari for all the sacrifices she has made so I can achieve my goals. She has given me a true sense of purpose since we met almost a decade ago and I have never been and never will be able to fully express my gratitude for all the joy, love and peace she has brought to my life.

Table of Contents

List of Figures

List of Tables

List of Appendices

Abstract

Incorporation of contextual inference in the optimality analysis of operational problems is a canonical characteristic of data-informed decision making that requires interdisciplinary research. In an attempt to achieve individualization in operations management, we design rigorous and yet practical mechanisms that boost efficiency, restrain uncertainty and elevate real-time decision making through integration of ideas from machine learning and operations research literature.

In our first study, we investigate the decision of whether to admit a patient to a critical care unit which is a crucial operational problem that has significant influence on both hospital performance and patient outcomes. Hospitals currently lack a methodology to selectively admit patients to these units in a way that patients individual health metrics can be incorporated while considering the hospitals operational constraints. We model the problem as a complex loss queueing network with a stochastic model of how long risk-stratified patients spend time in particular units and how they transition between units. A data-driven optimization methodology then approximates an optimal admission control policy for the network of units. While enforcing low levels of patient blocking, we optimize a monotonic dual-threshold admission policy. Our methodology captures utilization and accessibility in a network model of care pathways while supporting the personalized allocation of scarce care resources to the neediest patients. The interesting benefits of admission thresholds that vary by day of week are also examined.

In the second study, we analyze the efficiency of surgical unit operations in the era of big data. The accuracy of surgical case duration predictions is a crucial element in hospital operational performance. We propose a comprehensive methodology that incorporates both structured and unstructured data to generate individualized predictions regarding the overall distribution of surgery durations. Consequently, we investigate methods to incorporate such individualized predictions into operational decision-making. We introduce novel prescriptive models to address optimization under uncertainty in the fundamental surgery

appointment scheduling problem by utilizing the multi-dimensional data features available prior to the surgery. Electronic medical records systems provide detailed patient features that enable the prediction of individualized case time distributions; however, existing approaches in this context usually employ only limited, aggregate information, and do not take advantages of these detailed features. We show how the quantile regression forest, can be integrated into three common optimization formulations that capture the stochasticity in addressing this problem, including stochastic optimization, robust optimization and distributionally robust optimization.

In the last part of this dissertation, we provide the first study on online learning problems under stochastic constraints that are "soft", i.e., need to be satisfied with high likelihood. Under a Bayesian framework, we propose and analyze a scheme that provides statistical feasibility guarantees throughout the learning horizon, by using posterior Monte Carlo samples to form sampled constraints that generalize the scenario generation approach commonly used in chance-constrained programming. We demonstrate how our scheme can be integrated into Thompson sampling and illustrate it with an application in online advertisement.

Chapter 1.

Introduction

As we evolve through and beyond the information age, we transition from industrial mass production and generic service delivery mechanisms to individualized and knowledge-based operational systems. It is now a common practice for different services and industries to be more appreciative of data as they deploy personalized strategies and, therefore, almost every stage of any operation is a host to plenitude of data collected in various formats and magnitude. This trend has consequently motivated researchers to adapt and improve various machine learning and statistical inference methods to leverage big data. Thus, the literature is ripe with multitude of studies in predictive modeling focusing on retrieving certain trends in data to provide a better picture of the underlying system.

However, predictive models per se are not fit for operational decision making and optimality analysis simply because they are designed for a different purpose. Hence, to close this gap, a new line of research has emerged in the modern literature of analytics known as "Prescriptive Modeling" [20]. These techniques have equipped the decision makers with powerful machinery to take on routine operational problems in a more nuanced fashion such that the optimal operational decisions are made by absorbing the entirety or a faithful approximation of the available data [113]. In this dissertation we aim to design new

frameworks to build on prescriptive modeling and data-driven decision making literature.

The availability of Electronic Health Records (EHR) [4, 101] has positioned many of healthcare operations in the favorable realm of data-intensive problems. Similarly, with the expansion of world wide web, social networks, online cookies, etc. every interaction of an individual with the online world is traceable in high-dimensional data-sets and, therefore, accessible for optimizing digital marketing campaigns. Hence, for our operational cases, we focus on general problems faced by decision makers in healthcare and digital marketing.

This dissertation is presented in a multiple manuscript format. We provide three distinct frameworks to facilitate the integration of machine learning and optimization. In Chapter 2, we start by considering a separated estimation and optimization model to optimally allocate needier patients to a higher level of care based on their mortality risk and a comprehensive consideration of hospital operations. In Chapter 3, we discuss a new framework that enables the decision maker to directly incorporate data in performing optimization under uncertainty. For this chapter, we focus on the classic appointment scheduling problem. Finally, in Chapter 4, we transition to sequential optimization problems where the decision maker has to learn and decide as data becomes available sequentially. We focus on designing optimal strategies for digital marketing campaigns under probabilistic budget constraints.

In Chapter 2, considering the steep financial burden imposed by critical care units and the importance of patient outcomes, we investigate methodologies to improve admission decisions in the higher level of care units/wards. We generate a novel perspective to optimize a personalized risk-based admission policy. To do so, we leverage a recently developed estimate of mortality risk metric as a surrogate measure of each patient's true bed placement requirements. We integrate this metric with the hospital's bed capacity using a network model of patient flow through units/wards and a data-informed methodology. This study introduces a complex model of patient placement into a network of Intensive

Care Units (ICU) or Intermediate Care Units (IMC) and approximately optimizes the ICU and IMC patient admission policy to get the neediest patients into the ICU and IMC units using stochastic arrivals and patient flow processes. The results in Chapter 2 have appeared as individual research paper [98].

In the first part of Chapter 3, we study the surgical case duration problem. We offer a novel individualized perspective to address case duration prediction problem by leveraging the entire patient-hospital data available prior to the surgery. From the structured data we include patient related data, hospital and staff related data, operational data and temporal data. In addition to the structured data, to boost individualization, we also explore multiple text notes from surgeons and radiologists. Finally, we introduce generalizable predictive models and validate our results via extensive experiments on rich data sets of Memorial Sloan Kettering Cancer Center (MSKCC).

In the second part of Chapter 3, we investigate methods to incorporate individualized data-driven predictions into medical decision-making. As an application context, we investigate the surgery appointment scheduling problem, because the prediction of surgical case durations is a difficult and important challenge exhibiting high variability and uncertainty, particularly from the perspective of the individual patient. The existing approaches in this context usually employ only limited, aggregate information, and do not take advantages of the detailed available data. We show how the quantile regression forest (QRF) method, which provides an individualized predicted distribution, can be integrated into three common stochastic optimization approaches in addressing this problem, including stochastic optimization, robust optimization and distributionally robust optimization. We also numerically compare the benefits of our approaches relative to traditional models using a data from our Memorial Sloan Kettering Cancer Center.

Finally in Chapter 4, we provide the first study on online learning problems under

stochastic constraints that are "soft", i.e., need to be satisfied with high probability. These constraints are imposed on all or some stages of the time horizon so that the stage decision probabilistically satisfies a safety condition that is realized after the decision is made. The distributions that govern these conditions are learned through the collected observations. Under a Bayesian framework, we study a scheme that provides statistical feasibility guarantees through the time horizon, by using posterior Monte Carlo samples to form sampled constraints that leverage the scenario generation approach in chance-constrained programming. We demonstrate how our scheme can be integrated into Thompson sampling and illustrate it with an application in online advertisement using real data from Adobe Systems, Inc.

In the last Chapter of this dissertation we discuss the findings and provide concluding remarks. In addition, we provide future paths and potential research opportunities based on the outcomes of each chapter. Table 1.1 summarizes the main contributions in each chapter.

Table 1.1: Overview of the dissertation topics and their main contributions.

Topic	**Main Contributions**
Chapter 2: Data-driven Optimization Methodology for Admission Control in Critical Care Units	(1) Introducing a complex model of patient placement into a network of ICU and IMC units (2) Proposing a data-driven Mixed Integer Programming (MIP) optimization methodology in equilibrium that generates optimal key system metrics (3) Defining an individualized framework to selectively promote the highest risk patients to a higher level of care unit or demote patients of low risk to a lower level of care using the notion of mortality risk
Chapter 3: Optimization Under Uncertainty: Surgery Scheduling and The Promise of Individualization	(1) Proposing a state of the art predictive model for surgery duration estimation using the entirety of patient-hospital data (2) Defining an individualized stochastic optimization framework and showing the statistical consistency of this method (3) Introducing individualized robust optimization model, and showing the consistency and robustness guarantee of this technique (4) Proposing an individualized distributionally robust model for the surgery scheduling problem, and proving the statistical consistency of this approach
Chapter 4: Sequential Learning Under Probabilistic Constraints: Feasibility of Digital Marketing Campaigns	(1) Designing sequential methodologies that maintain probabilistic feasibility requirements with rigorous statistical guarantees (2) Providing the first study on a class of problems that not just include the rewards but also stochastic constraints that need to be satisfied with high probability (3) Illustrating the integration of our scheme into a variant of Thompson sampling

Chapter 2.

Data-driven Optimization Methodology for Admission Control in Critical Care Units

2.1. Introduction

Healthcare costs in the US are extremely high, with $759.1 billion annually in hospital costs alone [136]. Among critical care units/wards, Intensive Care Units (ICUs) are particularly expensive, with over $55 billion in estimated annual costs as of 2005 [63]. Although less expensive than ICUs, Intermediate Care units (IMCs) are also more costly than general/regular medical units. IMCs provide a level of care between intensive care and general units, and their inclusion in hospitals' network flow has shown benefits for both patients and system utilization [6]. These units have less equipment and a higher ratio of patients to nurses; therefore, IMCs are less costly to operate than ICUs. Considering the steep financial burden imposed by critical care units and the importance of patient outcomes, it is

essential to investigate methodologies to improve admission decisions in the higher level of care units/wards such as ICUs and IMCs. This research generates a powerful method and a novel perspective to optimize a personalized risk-based admission policy. We leverage a recently developed estimate of mortality risk metric as a surrogate measure of each patient's true bed placement requirements. We integrate this metric with the hospital's bed capacity using a network model of patient flow through units/wards and a data-informed methodology.

One important component of this problem is the management of a hospital units' bed capacity in the face of congestion. Hospital census (the number of patients in the hospital at a given time) and unit/ward level census are both variable, with a weekly pattern of variation (see [68] as one example in addition to others cited therein). Under-utilization of capacity reduces potential hospital revenue, whereas over-utilization leads to patient blocking, decreased quality of care, and job dissatisfaction among nurses [2, 68]. By blocking, we mean either that a patient is diverted to another hospital because the desired hospital is not accepting patients due to limited resources, or that a patient is placed in a different unit because the desired unit is full. Both types of blocking are undesirable. Placing a patient in a suboptimal unit may compromise quality of care and can lead to worse patient outcomes (i.e., increased length of stay, mortality rate, and readmission rate [40]).

The number of ICU beds in the U.S. is growing, but this growth is slower than the growth of demand for ICUs [64]. Some predict a shortage of ICU beds [139]. As the population in the U.S. continues to age rapidly, more stress will be placed on our healthcare systems because older patients require more frequent and more intensive healthcare services [126, 139]. Compounding this effect, hospitals across the country are reducing the number of available beds. This is especially troubling when we consider that the effects of ICU congestion are felt throughout the hospital, as patients are frequently competing for ICU

beds [78, 95]. This paper sheds light on how to efficiently allocate limited ICU and IMC resources.

ICU care is beneficial to patients [13, 122] and may decrease patient mortality rates during the first 24 hours of care, especially for those patients with high risk [122, 123]. ICU care is especially beneficial for critically ill patients; it has been shown to markedly improve outcomes for these critically ill patients [130]. Conversely, ICU care may not be beneficial for patients who are sufficiently well [130, 149].

Hence, some studies in the ICU admission literature include severity measures in the decision making processes [71, 121]. Such information may not be accessible until the patient is placed in the ICU, which limits its potential benefits. However, Cowen et al. have developed a mortality risk metric (MRM) [45] that is available prior to bed assignment. The MRM system has been operational in our partner hospital for a few years and we use this measure as one ICU admission criterion in our model, given the potential benefits of ICU care may be easier to demonstrate in patients with an elevated risk of dying [42]. Moreover, the MRM metric is not only associated with the risk of death, but also correlates significantly with the risk of other adverse health events [45]. The concept that drives implementation in our partner hospital is that MRM is an effective summary measure for a patient's comorbidities and her overall health state that complements the severity of the patient's chief diagnosis and acuity level for nursing care [44]. So, we are using MRM as a surrogate for a patient's need for an ICU/IMC bed to determine an optimal admission policy. But our method is general and could incorporate other metrics such as the Rothman index [112].

This paper introduces a more complex model of patient placement into a network of ICU and IMC units. Many patients are placed in the higher level of care units for reasons of nursing acuity, that is, the increased nursing hours per patient required relative to the

average patient, because of the complexities of the treatments administered and daily care needs. Because of the broader range of treatments available in ICU and IMC units, such patients require specific treatments available only in ICU or possibly in IMC unit. We call patients who are either high acuity, or simply require a treatment only found in an ICU or IMC, "mandatory" patients (and we differentiate ICU mandatory from IMC mandatory to designate the unit that they need) [93]. In contrast to mandatory placements in ICUs and IMCs, we consider a third placement criterion, risk of dying, that may qualify a patient for a higher level of care. We call patients who are not mandatory, i.e., who do not require treatment specific to an ICU or IMC unit, "discretionary". This distinction allows us to investigate the opportunities to admit critical discretionary patients to ICUs or IMCs. This is expected to result in improved health outcomes and reduced mortality. In essence, we study the trade-off between blocking future mandatory patients and admitting high risk discretionary patients.

Let us consider some of the issues that motivate optimizing the admission control to ICUs and IMCs. Currently, patients are regularly placed in the ICU who could receive equally beneficial care in a non-ICU unit. For instance, most low risk patients never use specific ICU services when in the ICU (low risk patients are here defined as those with mortality risk $< 0.06\%$) [109, 149]. This seems to be present in the data from our partner hospital. During an eighteen-month period the number of ICU admissions for low risk patients were 20% higher than high risk patients while a considerable number of high risk patients were still admitted to the non-ICU units. Some hospitals use Intermediate Care Units in addition to ICUs to provide an intermediate level of care. Although less than for ICUs, there is significant competition for IMC beds, which are much scarcer than regular/general beds. Given the limited resources available in the critical care units, this motivates the necessity of a methodological framework to admit patients to higher levels

of care in a selective manner.

There is a need in the medical community for a good predictive model identifying which patients would and would not benefit from ICU or IMC care. Lacking good information, doctors ration (i.e., deny an ICU bed to some patients to manage capacity) less accurately than they would given perfect information. The current practice of assigning ICU beds is often 'subconscious rationing' rather than a data driven policy [63]. Doctors make resource allocation decisions based on clinical judgment at the moment, which may or may not be cost effective or beneficial in the long run [126]. Additionally, some studies suggest that doctors are overconfident in their ability to judge ICU need [126]. Given this state of practice, an approach to ICU admissions that uses scientific tools and methods seems beneficial. It is worth noting that although our partner hospital has a "triage" coordinator overseeing admission to bed units, the hospital has still pursued the use of a personalized mortality risk metric as a key component of its admission decision process. With the advent of electronic medical records (EMRs), new possibilities are emerging for data driven, analytical models that use past data to inform future behavior.

The main thesis of our study is that we can approximately optimize the ICU and IMC patient admission policy to get the neediest patients into the ICU and IMC units. This is expected to improve patient outcomes by reducing patient complications and mortality as well as reducing unplanned transfers to the ICU. Reducing transfers is important because of the longer length of stays and higher mortality rate observed in transferred patients [57]. Unplanned transfers from a non-ICU bed to the ICU are especially harmful because in this case the patient needs immediate transfer or her condition is likely to deteriorate rapidly, and an appropriate bed cannot always be found sufficiently quickly [116]. In addition, considering stochastic arrivals of specific patient types and their historical paths in the hospital, we will be able to adjust the admission decision such that it delivers patient

probabilistic needs while tracking units' utilization to support hospital profitability.

Unlike the prior literature, we consider an integrated model of several ICUs and several IMCs, and we aggregate all the other "regular" bed units of the hospital in a unit we call General Medical Beds (GMB). This allows us to focus on the flow details of the higher level of care units while accounting for total bed capacity of the hospital and the fact that patients in the GMB units may at any time during their stay require elevation of care to an IMC or ICU bed.

Built upon approximations of the stochastic arrival and flow processes, this paper contributes a data-driven Mixed Integer Programming (MIP) optimization methodology in equilibrium that can generate optimal key system metrics. Different policies can be applied to enforce desired managerial perspectives. We will consider two admission strategies. The first strategy requires the model to generate constant admission thresholds for every day of the week, while the second strategy allows daily variation across the days of the week. We show that the latter will not only increase the number of admitted patients to critical units, but it will also reduce the blocking rates in the long run. We validate our results via extensive numerical simulations on synthetic and real data sets.

2.2. Literature Review

Green showed that as many as 90% of hospitals lack capacity to accommodate high risk patients arriving into their ICUs [62]. To address this issue, there has been a rise in ICU management literature in recent years. These studies can be classified into two broad categories: 1) ICU capacity management and 2) ICU admission control. Chalfin et al. found that in case of full capacity, new ICU admission requests will be delayed and patients with critical conditions may be boarded in the emergency department [37]. This can endanger the patient's life as delays in ICU admission can significantly raise the mortality

rate [22, 37].

On the admission side, many researchers employed queueing theory and Markov decision process theory to study the admission control and elective admission policies for critical care units [120, 145]. These studies focus on a single patient type and/or a single resource, and disregard hospital network effect. Kim et. al. examined the ICU admission problem via econometrics [79]. They showed the importance of ICU admission policies and derive a threshold policy for admitting patients to the ICU. While Kim et al. studied a single unit, this paper models multiple patient classes and a network of units in the problem of ICU admission policies.In this paper we clarify that many patients require one or more specific treatments that are available only in an ICU of a specific type or in some cases a set of IMC units. Kim et al. assume that patients can go to any ICU, regardless of diagnosis, so they only consider one combined ICU. Our collaboration with hospital staff, indicated that, in our partner hospital, patients are rarely placed off-service (i.e., in the wrong type of ICU), and if they are, it is only for a few hours until a bed opens up in the desired ICU. In this paper, our optimization based model will examine the ICU and IMC admission problem when considering multiple ICUs and IMCs, where different units treat different types of patients.

Mathews and Long also study the ICU admission problem [93]. Their contribution builds on the work of KC and Terwiesch [78] by considering the impact of census on patient non-service wait time in the ICU, namely how long patients wait to be transferred after they are done with treatment, as a function of census. Although they consider the emergency department (ED) versus non-ED distinction, acute and non-acute conditions for each patient may not be easily interpretable and measurable in practice. We model the class of mandatory ICU or IMC placed patients who cannot obtain a required treatment in a GMB type bed. This has a large impact on the accuracy of the results of the models.

For instance, mandatory patients comprise almost 20% of the overall body of patients in our partner hospital.

Step Down Units (SDUs) are considered to be reliable alternatives for ICUs, and there has been a surge in studying the influence of these units on patient outcome [39] and hospital performance [6]. Mathews and Long consider both an ICU and an IMC for a single type platform. In contrast, we will examine the ICU admission problem with multiple ICUs and IMCs. This is important for two main reasons. First, a patient's diagnosis often uniquely determines which ICU or IMC the patient should go to. It can be a patient safety issue if a patient is placed in the wrong ICU or IMC. Second, there are complex network flows between the various ICUs and IMCs, which affect overall system performance. We should clarify that IMCs are very similar to Progressive Care Units (PCUs) or Step Down Units. IMCs are nearly identical in character to SDUs, and the main difference between IMCs and PCUs is that PCUs have more flexible staffing, allowing a patient's nursing ratio to shift as her care needs change. Whereas a patient would be transferred out of an IMC when her care needs lessen, in a similar situation she would remain in a PCU, just with a lower nurse:patient ratio. This distinction is not essential to this paper, and though we will talk about IMCs for the remainder of the paper, one can apply our results to PCUs or SDUs as staffing decisions are not our focus.

Recently, Helm, Van Oyen, and coauthors developed methodologies for optimal elective hospital admissions [67, 68]. Of these, [68] is by far the closest to this study. We extend these methods to the critical care environment by introducing a novel individualized/personalized framework. The primary focus of elective admission models exploits the ability of elective patients to be delayed until a later day that would improve the patient flow. This can be understood as a change in the weekly scheduling of surgeries to achieve a weekly block-scheduling pattern that facilitates patient flow. Our paper takes patients

as they arrive and selectively promotes the highest risk patients to a higher level of care unit or demotes patients of low risk to a lower level of care. This purpose requires a new definition of patient types, the approach to the modeling of patient flow pathways and the complex math programming formulation. We also develop an exact approach for computing the standard deviation of the census in our model. Interestingly, this paper develops a methodology to compute an optimal dual-threshold admission policy for IMCs and ICUs.

2.3. System Design and Modeling Assumptions

We generate optimal admission policies through a binary matrix, $[\theta^{\mathbf{k}}_{\mathbf{udi}}]$, whose elements, $\theta^{\mathbf{k}}_{\mathbf{udi}}$, are the risk threshold decision variables of our MIP model indicating whether we admit a patient of class k with risk level i to unit u on day d. Once the binary matrix is constructed, we can then create a daily dual-threshold admission policy. Individual ICU and IMC admission policies will be easily accessible by tracking decision variables in the matrix. We check the monotonic admission-risk patterns and define the threshold for ICUs, IMCs and general beds as a global rule. To establish such well-defined thresholds, however, we need to discretize the mortality risk metric into several bins such that we have roughly equally frequent observations in each risk interval. For instance, if we consider r different risk levels, then the admission rule to unit u for patients of class k on day d is fully defined by the array $[\theta^k_{ud1}, \theta^k_{ud2}, ..., \theta^k_{ud(r-1)}, \theta^k_{udr}]$.

To further differentiate mandatory patients, we will call the group of patients who require treatment that is only available in the ICU, ICU-mandatory. These ICU-mandatory patients should receive priority for ICU admission whenever there is a conflict with a discretionary patient. Similarly, the IMC-mandatory patients should receive priority for IMC admission. Although we usually treat them separately and attach a desired unit for each patient, when convenient we will sometimes lump ICU-mandatory and IMC-mandatory

patients as a single group of mandatory patients.

In our optimal admission policy for critical care units, we consider an infinite horizon and cyclo-stationary model with a seven-day week as the period. Thus, in equilibrium, the same policy is applied every week, but it may vary by day of week to accommodate day of week effects. This model matches the weekly schedule of clinics, operating rooms, etc.

We work on a dataset that spans two years, with more than $200,000$ data points representing over $70,000$ distinct patients at our partner hospital. The dataset contains the information regarding patients' mortality risk, specialty, arrival and discharge, and the units they visited during their stay by date and time. The hospital network consists of three different ICUs as well as five distinct IMCs.

For the sake of intuition, before developing the more realistic model, we employ a simple queueing model with a single patient class to provide insight into how the admission to ICU's based on a mortality risk threshold reveals the average sum of MRM's as a function of the threshold and the number of beds. This simple single-class $M/M/s/s$ model with rejection when an arrival finds the ICU full with s patients assumes Poisson arrivals and exponentially distributed length of stay (LOS) in the ICU. We will revisit these assumptions in Section 2.3.1.

Our scoring metric is the time-average sum of MRMs over patients in an ICU bed. We will discuss this objective function in more detail below and extend it to consider IMC beds also in the network model. The idea is that the patient mortality outcomes are proportional to the mortality risk, which justifies maximizing the average mortality in these beds that best care for the sickest patients within the bed capacity. The relationship between mortality risk and poor health outcomes is not known precisely; however, the mortality risk is associated with the risk of unplanned transfers, cardiopulmonary arrests, intra-hospital complications and 30 day readmissions [45]. Therefore, we believe this objective is appro-

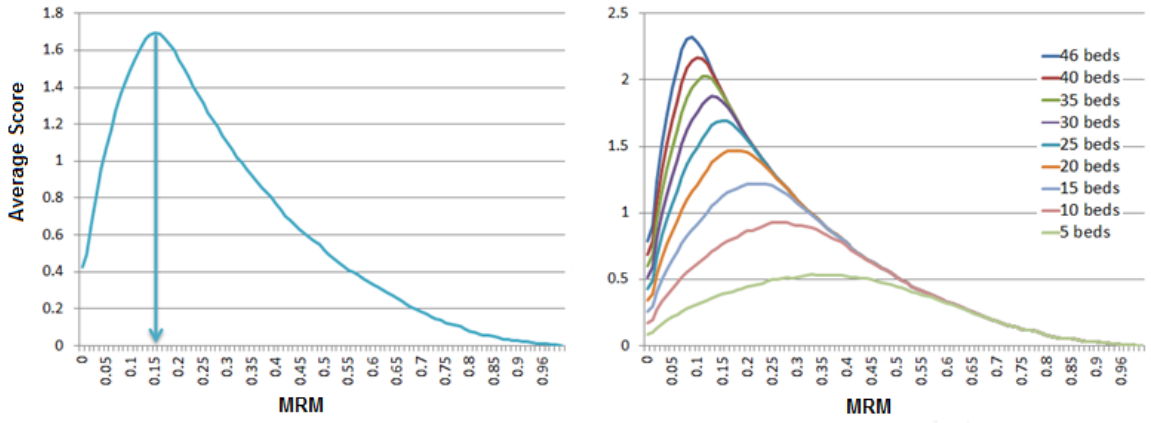

Figure 2.1: Left: Score as a function of the MRM threshold used, when the ICU has 25 beds. The arrow indicates the optimal threshold. Right: Score as a function of MRM threshold, for 5 (bottom curve), 10, 15, 20, 25, 30, 35, 40, and 46 (top curve) beds in the ICU.

priate within the current limits of knowledge. Figure 2.1 indicates that for a given demand model, as the ICU capacity increases the optimal MRM threshold decreases. It also shows that as we reduce the MRM threshold below the optimal level, the average score declines, which means that initial admission of lower risk patients impeded the entrance of more critical patients. Similarly, raising the MRM threshold above the optimal level diminishes the overall score of the system, which reflects the underutilization of ICUs.

2.3.1. From Queueing to Optimization

The above queueing model gives good intuition about the dynamics of the system, and it effectively communicates the concept to the hospital staff. On the other hand, the model is limited by the over-simplification of the problem. Therefore, in order to gain a more accurate and general framework, as well as to get richer results, we create a Mixed Integer Programming (MIP) model that allows us to build a data driven model and to capture network effects, including the interplay between the ICUs, IMCs, and the rest of the hospital. For instance, if a patient is admitted to a specific ICU on a given day, there

is a chance (which depends on the patient's characteristics) that she will be in a specific IMC after a given number of days. See Figure 2.2 for an illustration of possible patients' movements in the hospital network, which allows for multiple visits to a station (feedback), and extends from the time of admission until discharge from the hospital. We will discuss in the next section how this stochastic patient stay pattern is realized in our modeling framework.

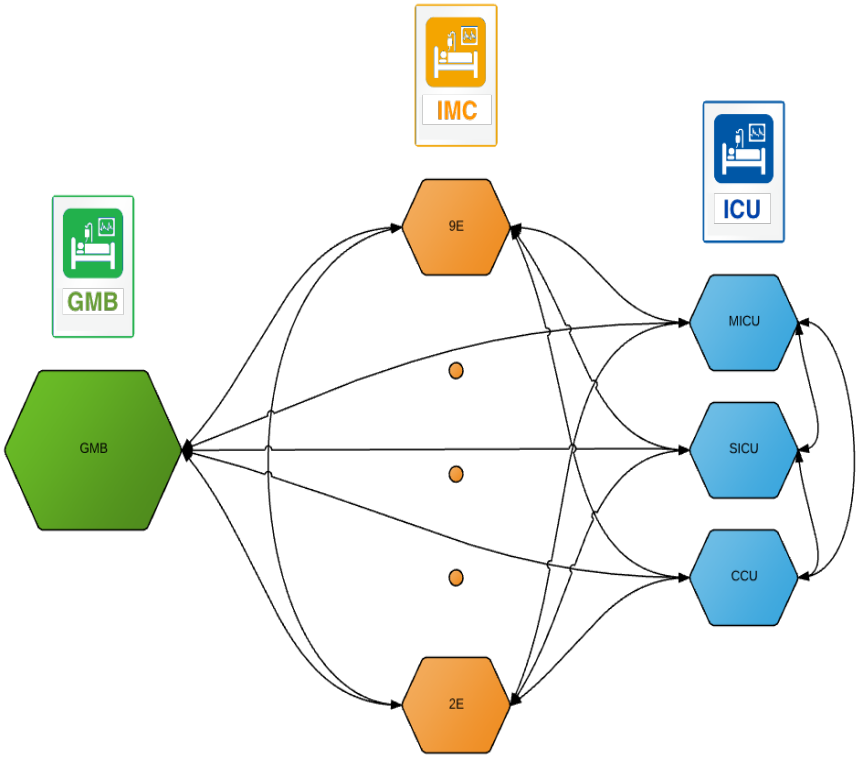

Figure 2.2: An illustration of the network of ICUs, IMCs, and the rest of the hospital. We have aggregated the rest of the hospital into one single general medical bed unit (GMB).

Employing an MIP model, we can find an optimal dual-threshold-type policy. The dual-threshold policy is reasonable and probably optimal if one accepts the preferences of the clinical leadership. They specifically desire to announce the MRM risk threshold for ICU admission and the lower one for IMC admission. A dynamic, state-dependent policy

should perform better; however, a costly and complex decision support system would be required. A patient with an MRM score at or above the ICU-threshold, is approved for ICU admission, and will then be considered for IMC admission if there is no room in the ICU. If patient does not meet this threshold but meets the IMC-threshold, the patient is approved for IMC admission. A patient who does not meet either of these thresholds is sent to a non-ICU/IMC unit, which is modeled using the GMB unit to avoid unnecessary complications in the model. Our partner hospital is already using MRM thresholds to guide ICU and IMC admission decisions, but the thresholds were not optimized prior to this work.

We note that an MIP framework enables us to vary the thresholds by day of week and by unit. Another advantage of the MIP modeling is that we can introduce multiple classes of patients, which is crucial for our problem as we are dealing with mandatory and discretionary patients. This methodology, therefore, lets us use a threshold policy for discretionary patients while automatically approving all mandatory patients for ICU or IMC admission. We can also approximate the queueing dynamics.

2.3.2. Mixed Integer Program Formulation

Our goal is to compute an admission policy that will admit the neediest discretionary patients into the ICUs and IMCs. To do this, we propose the following: mandatory patients (i.e., those who require treatment specific to the ICU or IMC) are always approved for admission to the relevant unit. To be specific, ICU-mandatory patients are always approved for admission to the appropriate ICU and similarly for IMC-mandatory patients. All other patients, who we are calling discretionary, may be approved for admission to an ICU or IMC based on their respective MRM scores. We will apply the dual threshold-type policy on these discretionary patients with separate thresholds for ICU admission and for

IMC admission. The intuition here is that high risk patients would benefit most from ICU care, but would also benefit from IMC care, and that less risky patients would benefit from IMC care without requiring the extra level of care given by ICU admittance.

We explained briefly in the introduction why our partner hospital believes the degree of a patient's neediness can be approximately quantified by her MRM risk score and our objective is to maximize the sum of risk scores over patients admitted to the ICUs and IMCs. Based on discussions with the medical staff at our partner hospital, we found that ICU and IMC admissions are roughly equally beneficial for discretionary patients. Thus, we will have a zero reward for placement in the aggregate GMB unit and a reward equal to the MRM for admitting a patient to the ICU or IMC; however, this only makes sense in light of a constraint that the risk threshold for ICU admission is strictly higher than that for IMC admission. If one had data to support it, one could easily introduce a parameter to proportionally weight the reward for admitting a patient to various units. We are interested in an optimized, capacity-based admission policy, rather than myopic acceptance which would be suboptimal. To capture this, we multiply each discretionary patient's expected risk score by an indicator variable for patient admission to an ICU or IMC unit. This can be thought of as integrating, over the course of the study horizon, the risk score of each patient whose initial unit of admission was any of the various ICUs and IMCs. Thus, our objective function is to maximize

$$\sum_{d\in\mathcal{D}}\sum_{k\in\mathcal{K}}\sum_{i\in\mathcal{I}\setminus\{\mathcal{M},\mathcal{M}'\}}\sum_{u\in\mathcal{U}\setminus\{\text{GMB}\}}\theta^{\mathbf{k}}_{\mathbf{udi}}R_i\ , \tag{2.1}$$

where $\mathcal{I}$ is the set of risk groups (obtained by partitioning patients by MRM) and $\mathcal{M}$ and $\mathcal{M}'$ define the IMC and ICU mandatory patients, respectively. Also, $\mathcal{D}$ is the set of days of the week (Sunday, Monday, etc.), $\mathcal{U}$ is the set of units/wards, and $\mathcal{K}$ is the set of

patient classes. In this context, a patient's class is determined by her treatment needs and diagnosis; we will use this to determine to which units a patient may be admitted to. For instance, a cardiac class patient would likely go to the cardiac ICU instead of the surgical ICU (SICU) if she required ICU services. Our decision variable, $\theta^{\mathbf{k}}_{\mathbf{udi}}$, is a binary variable that is equal to 1 if we allow patients of class k and risk group i to be admitted to unit u on day d, and equal to 0 otherwise. We let $\mathbf{\Theta}$ refer to the matrix $[\theta^{\mathbf{k}}_{\mathbf{udi}}]$. Though we do not have an explicit threshold variable, the thresholds can be inferred from $\mathbf{\Theta}$. R_i is the expected risk score for a patient in risk group i. Our objective is simple and fair, admitting as many of the stratified high-risk groups of patients as capacity will allow and admitting higher risk patients over lower risk ones.

To link the operational constraints of the problem to our objective function, we have to characterize the stochastic paths of different patient classes who are admitted to ICUs, IMCs, and the regular beds of GMB. Incorporating binary decision variables, we can define a new version of the stochastic location process introduced by Helm et. al. [68]. In that work the same patient population was admitted to the same unit, but the day of week of admission could be changed to improve the patient flow. Emergency patients were not changed, but elective patients could be delayed until a day of admission that would improved the flow of the hospital. That paper also considered a version that increased the hospital throughput by selectively admitting more patient of any type desired subject to the capacity limits of the hospital.

In our setting, the patient of a given type cannot wait to be admitted. Moreover, we intentionally selectively change the first unit of admission based on his/her personalized risk metric. Some discretionary patients will be elevated to a higher level unit (and a distribution on the initial unit of admission must be used to capture detailed patient type information that is not explicitly modeled). Other discretionary patients with a low risk

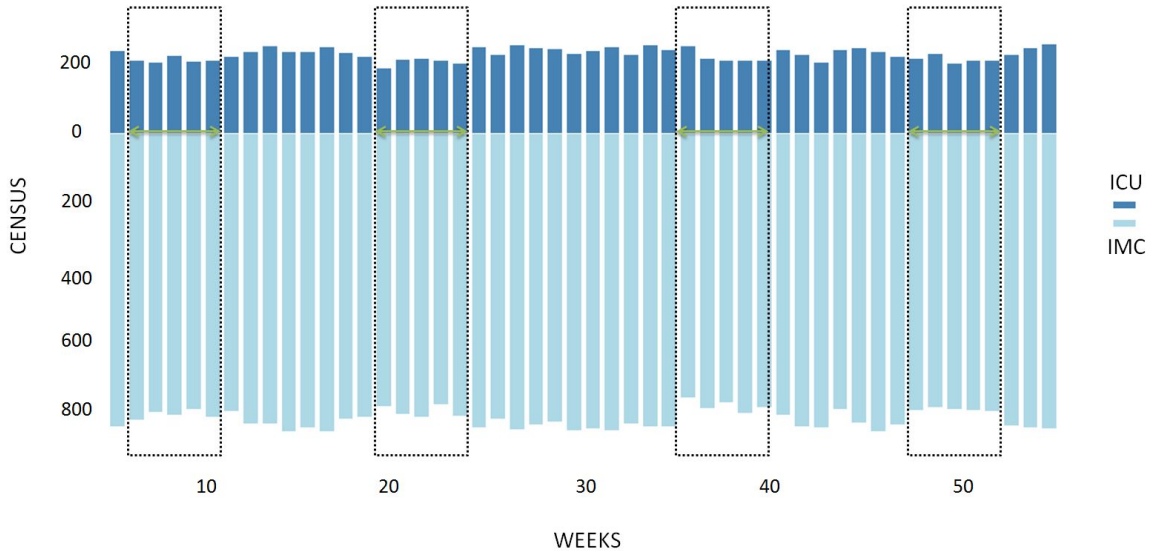

Figure 2.3: Identifying less congested periods in weekly census analysis of ICUs and IMCs, shown in the boxes

will be demoted to a lower level of care. Including a patient's mortality risk conveys a broad range of information related to a patient's medical history, lab results, and demographics [45]. This will improve the model's performance in replicating patient flow in the most appropriate units of the network.

We must construct a model of how long patients will spend in any given unit and their likelihood of moving to another unit. We view this as a stochastic location process over time, which is indexed to the weekday of admission and the patient's class. This can be defined as a stochastic model that captures length of stay (in days) in each unit and the units that patients visit during their hospital stay. Additionally, we assume that the stochastic itineraries of different patient classes are independent throughout the hospital, which fits with our offered load approach to the census distribution. Also, to derive the stochastic location of patients, we build the model based on the less congested periods in our dataset to preserve the independent itinerary assumption (Figure 2.3).

Without any constraints, it is easy to see that the optimal admission policy would

prescribe $\theta^{\mathbf{k}}_{\mathbf{udi}} = 1$ for all k, u, d and i, meaning that everyone is approved for ICU and IMC admission. This corresponds to the intuition that, given infinite resources, it would be beneficial to make the entire hospital into an ICU. No hospital can afford that expense; however, so we need to constrain the system by the available capacity while maintaining a certain service level. Of course, one can use this approach to quantify the higher level of care capacity needed to send all patients above a given risk level to an ICU or similar objective.

To formally define our stochastic location process, we will employ a modified version of the stochastic location process modeling approach developed in [68]. Let $P^{u'}_{d'kui}(r)$ be the probability that a patient of class k and risk group i who arrived on day d' is in unit u, r days after admission, given that the patient was initially admitted to unit u'. This measure is evaluated on a weekly basis for n weeks ($r = d - d' + 7n$). Also, let $\bar{F}_{d'ki}(a)$ be the probability that there are more than a arrivals of class k and risk group i on day d' and consider n as the counter of the weeks.

To compute the congestion under any policy, given $\mathbf{\Theta}$, we seek estimates of the mean and variance of the census for each unit on each day. As will be described below, using the mean and variance of a stochastic location process, we can approximate the units' capacity and overall patient flow in the hospital. By keeping a sufficiently low probability that the census deviations above the mean go beyond the capacity of the unit, we can limit the blocking and thereby ensure that the new system will work within the hospital's unit network capacity constraints. Small amounts of blocking are not a practical problem, because the hospital is very experienced at reacting to blocking. The following theorems show that the mean and the variance of the units' daily census can be calculated linearly in $\mathbf{\Theta}$.

Theorem 2.1. *The mean census in unit u on day d is given by*

$$\mu_{\mathbf{du}} = \sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty}\sum_{a=0}^{\infty} P^{u'}_{d'kui}(r)\bar{F}_{d'ki}(a)\theta^{\mathbf{k}}_{\mathbf{u'd'i}}.$$

Theorem 2.2. *The variance of census in unit u on day d is given by*

$$\mathbf{v_{du}} = \sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty}\sum_{a=0}^{\infty}\{P^{u'}_{d'kui}(r)\bar{F}_{d'ki}(a) - [P^{u'}_{d'kui}(r)\bar{F}_{d'ki}(a)]^2\}\theta^{\mathbf{k}}_{\mathbf{u'd'i}}.$$

Theorems 2.1 and 2.2 state that the mean and variance of daily unit census can be calculated linearly in $\mathbf{\Theta}$ for any arrival process. Clearly, these are data-driven estimations based on historical records for each risk stratum/group, which were available by limiting the number of risk strata in the model without assuming a specific arrival distribution. The proofs for these theorems are given in Appendix A.2.

We incorporate finite capacity by limiting blocking. By blocking, we mean that a patient cannot be admitted to the preferred unit because it is full. If we do not constrain blocking, our model would lead to a first-come-first-served policy, which would be unacceptable due to the blocking of many of the incoming mandatory patients. To enforce a required service level for our system, we model the right tail of the census of each unit on each day of the week using its own normal/Gaussian distribution. Note that the mean and variance of the census are calculated by Theorems 2.1 and 2.2, respectively. The normality of workloads is common in healthcare workload analysis literature, and it has been shown to be a good model in similar settings such as [72]. Admittedly, allowing for non-normal distributions considerably complicates the model and often leads to nonlinearity of the optimization. Because stability requires that the capacity must exceed average demand for service, the

key is to estimate the overload in each unit by examining the right tail of the census distribution, given a desired service level α. This is captured in the equation below:

$$\mu_{\mathbf{du}} + z_\alpha \sigma_{\mathbf{du}} \leq C_{du} \tag{2.2}$$

where $\mu_{\mathbf{du}}$ is the mean census in unit u on day d given $\mathbf{\Theta}$, $\sigma_{\mathbf{du}}$ is the standard deviation of the census in unit u on day d given $\mathbf{\Theta}$, z_α is the standard normal z-score coefficient that corresponds to service level α, and C_{du} is the capacity of unit u on day d. C_{du} may vary by day of week because some units may have different staffing levels on the various days of the week, leading to different effective unit capacities for those days. The equation above implies that, on day d and in unit u, patient blocking will take place with probability less than or equal to $1 - \alpha$.

Now that we have a mechanism to capture patients' blocking probabilities in our system, we can set the α such that the desired expected number of blockings by unit and day of week is respected. The optimization will constrain the blocking probability $(1 - \alpha)$ so that the desired expected number of blockings by unit and day of week is kept to a small fraction of arrival volume, but this can be modified and applied for different hospitals in accordance with their policies and requirements.

In the derivation of $\mathbf{v_{du}}$, the variance of the census, we find that it is linear in $\mathbf{\Theta}$ (Appendix A.2). This would make the standard deviation non-linear in $\mathbf{\Theta}$, as $\sigma_{\mathbf{du}}$ is the square root of $\mathbf{v_{du}}$, which in turn would make (2.2) nonlinear. To avoid this issue, we achieve linearity by first rearranging (2.2) as

$$z_\alpha \sigma_{\mathbf{du}} \leq C_{du} - \mu_{\mathbf{du}} \ . \tag{2.3}$$

We, then, raise both sides of (2.3) to the second power to replace $\sigma_{\mathbf{du}}$ with $\mathbf{v_{du}}$:

$$z_\alpha^2 \mathbf{v_{du}} \leq C_{du}^2 + \boldsymbol{\mu}_{\mathbf{du}}^{\mathbf{2}} - 2C_{du}\mu_{\mathbf{du}} \ . \tag{2.4}$$

The equation above is equivalent to (2.3) only when both sides of (2.3) are non-negative. Hence, we enforce $C_{du} - \mu_{\mathbf{du}} \geq 0$ as a new set of constraints in our model. To linearize this equation, we take advantage of the binary characteristics of our main decision variables, $\theta_{\mathbf{udi}}^{\mathbf{k}}$. As defined in theorem 2.1, μ is constructed via a linear combination of $\theta_{\mathbf{udi}}^{\mathbf{k}}$. Therefore, given $(\theta_{\mathbf{udi}}^{\mathbf{k}})^2 = \theta_{\mathbf{udi}}^{\mathbf{k}}$, the square of mean census can be written as:

$$\begin{aligned}\boldsymbol{\mu}_{\mathbf{du}}^{\mathbf{2}} = &\sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty}\sum_{a=0}^{\infty}\{P_{d'kui}^{u'}(r)\bar{F}_{d'ki}(a)\}^2\theta_{\mathbf{u'd'i}}^{\mathbf{k}} \\ &+\sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{\substack{i'\in\mathcal{I}\\ i'\neq i}}\sum_{\substack{d''\in\mathcal{D}\\ d''\neq d'}}\sum_{\substack{u''\in\mathcal{U}\\ u''\neq u'}}\sum_{\substack{k'\in\mathcal{K}\\ k'\neq k}}\theta_{\mathbf{u'd'i}}^{\mathbf{k}}\theta_{\mathbf{u''d''i'}}^{\mathbf{k'}}\end{aligned} \tag{2.5}$$

It can be seen that (2.5) includes multiplications of two binary decision variable, which results in nonlinearity. However, we can avoid this problem by introducing new binary variables and a new set of constraints. In fact, we can replace each of the nonlinear terms generated in (2.5) by one new binary variable and three new constraints. Let's call our new binary variables ω. Now, the linear equivalent of (2.5) can be defined by the following equations:

$$\begin{aligned}\boldsymbol{\mu}_{\mathbf{du}}^{\mathbf{2}} = &\sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty}\sum_{a=0}^{\infty}\{P_{d'kui}^{u'}(r)\bar{F}_{d'ki}(a)\}^2\theta_{\mathbf{u'd'i}}^{\mathbf{k}} \\ &+\sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{\substack{i'\in\mathcal{I}\\ i'\neq i}}\sum_{\substack{d''\in\mathcal{D}\\ d''\neq d'}}\sum_{\substack{u''\in\mathcal{U}\\ u''\neq u'}}\sum_{\substack{k'\in\mathcal{K}\\ k'\neq k}}\omega_{\mathbf{u'u''d'd''ii'}}^{\mathbf{kk'}}\end{aligned} \tag{2.6}$$

$$\omega^{\mathbf{kk'}}_{\mathbf{u'u''d'd''ii'}} \leq \theta^{\mathbf{k}}_{\mathbf{u'd'i}} \tag{2.7}$$

$$\omega^{\mathbf{kk'}}_{\mathbf{u'u''d'd''ii'}} \leq \theta^{\mathbf{k'}}_{\mathbf{u''d''i'}} \tag{2.8}$$

$$\omega^{\mathbf{kk'}}_{\mathbf{u'u''d'd''ii'}} \geq \theta^{\mathbf{k}}_{\mathbf{u'd'i}} + \theta^{\mathbf{k'}}_{\mathbf{u''d''i'}} - 1 \tag{2.9}$$

Hence, we are able to overcome the nonlinearity of the workload control constraint by using (2.6) in (2.4). It is worth mentioning that the nonlinearity issue imposed by the standard deviation of unit census in (2.3) could be alternatively resolved by applying numerical approximation. One can also apply Newton's method to approximate $\sigma_{\mathbf{du}}$ by $\frac{1}{2}\left(\frac{\mathbf{v_{du}}}{\hat{\sigma}_{du}} + \hat{\sigma}_{du}\right)$, where $\hat{\sigma}_{du}$ is the historical standard deviation of the census in unit u on day d. Plugging it into our initial equation (2.3) gives:

$$\mu_{\mathbf{du}} + \frac{z_\alpha}{2}\left(\frac{\mathbf{v_{du}}}{\hat{\sigma}_{du}} + \hat{\sigma}_{du}\right) \leq C_{du} \tag{2.10}$$

This approximation can avoid the use of extra binary decision variables and constraints to mitigate the complexity of the algorithm. We introduce the above exact methodology; however, the approximate formulation of the problem can also be found in the Appendix A.1.

In addition to the operational constraints discussed above there are other technical and problem-specific considerations that we need to observe. As mentioned earlier, we have to ensure that all ICU-mandatory and IMC-mandatory patients are approved for ICU and IMC admission, respectively. We do so by requiring:

$$\sum_{u \in \mathcal{U}^{ICU}} \theta^{\mathbf{k}}_{\mathbf{udM}} = 1, \tag{2.11}$$

$$\sum_{u \in \mathcal{U}^{IMC}} \theta^{\mathbf{k}}_{\mathbf{udM'}} = 1, \tag{2.12}$$

where superscript $i = M$ denotes that a patient is ICU-mandatory, $i = M'$ denotes that a patient is IMC-mandatory, $\mathcal{U}^{ICU}$ is the set of ICUs ($\mathcal{U}^{ICU} \subset \mathcal{U}$), and $\mathcal{U}^{IMC}$ is the set of IMCs ($\mathcal{U}^{IMC} \subset \mathcal{U}$).

We also ensure that each patient stream is initially sent to only one unit. We can partition the patient streams in such a way that splitting patient streams between units would give no advantage. It is worth noting that we are aggregating the main hospital (i.e., all GMB units) into one unit because, for our purposes, there is little value in modeling those flows in detail.

$$\sum_{u \in \mathcal{U}} \theta^{\mathbf{k}}_{\mathbf{udi}} = 1 \tag{2.13}$$

Our model must also respect limitations beyond units' capacities and the mandatory decision rules. Therefore, we constrain the number of admitted patients of a certain class to each unit on each day to be less than the maximum number allowed by the hospital:

$$\lambda^{k}_{di} \theta^{\mathbf{k}}_{\mathbf{udi}} \leq \beta^{\mathbf{k}}_{\mathbf{udi}} \ , \tag{2.14}$$

where λ^{k}_{di} is the mean arrival rate of patients of class k and risk group i on day d, and $\beta^{\mathbf{k}}_{\mathbf{udi}}$ is the maximum allowed number of admissions for patients of a certain class with a certain risk level to each unit on each day. $\beta^{\mathbf{k}}_{\mathbf{udi}}$ can be easily set to the maximum historical value from the data. This constraint also ensures that the initial admission unit is conforming to the patient requirements and the current hospital practice.

It is also important to note that we want to maintain threshold level consistency for all days and in all units such that on a given day no higher risk patient goes to a lower level of care while a lower risk patient is admitted to a higher level of care. To ensure this pattern,

we have to include monotonicity constraints that enforce the required structure:

$$\theta^{\mathbf{k}}_{\mathbf{udi}} \leq \theta^{\mathbf{k}}_{\mathbf{ud(i+1)}} \qquad \forall u \in \mathcal{U}^{ICU}, \tag{2.15}$$

$$\theta^{\mathbf{k}}_{\mathbf{udi}} \leq \theta^{\mathbf{k}}_{\mathbf{u'd(i+1)}} \qquad \forall u \in \mathcal{U}^{IMC}, u' \in \mathcal{U}^{ICU}, \tag{2.16}$$

$$\theta^{\mathbf{k}}_{\mathbf{udi}} \leq \theta^{\mathbf{k}}_{\mathbf{u'd(i+1)}} \qquad \forall u \in \mathcal{U}^{GMB}, u' \in \mathcal{U}^{IMC}, \tag{2.17}$$

where $\mathcal{U}^{GMB}$ refers to general medical beds. Please also note that the risk groups are set to be ordinal. Additional constraints can be added, if desired. For instance, one could enforce a constant threshold across units and/or days of the week.

2.4. Numerical Results

In this section we validate and compare our results using almost two years of data from our partner hospital corresponding to over $70,000$ distinct patients. We first validate the mean census process estimations for different units on different days of the week. This hospital has three ICUs and five IMCs, and we labeled all other units as general medical beds. We confirm that our estimation of the unit workload nearly equals the actual unit census throughout the hospital by day of week.

Table 2.1 compares the weekly mean census for the estimated model versus the historical observation. The results indicate that the errors of our estimates derived from Theorems 2.1 and 2.2 are relatively negligible. The results also match closely with respect to different days of a week on a unit-by-unit basis. Figure 2.4 illustrates that our model performs very well in estimating both mean and standard deviation of the census for one of the ICUs in our network (MICU).

Having a reliable census process approximation, we can perform optimization to deliver operational measures such as admission thresholds given a desirable service level without

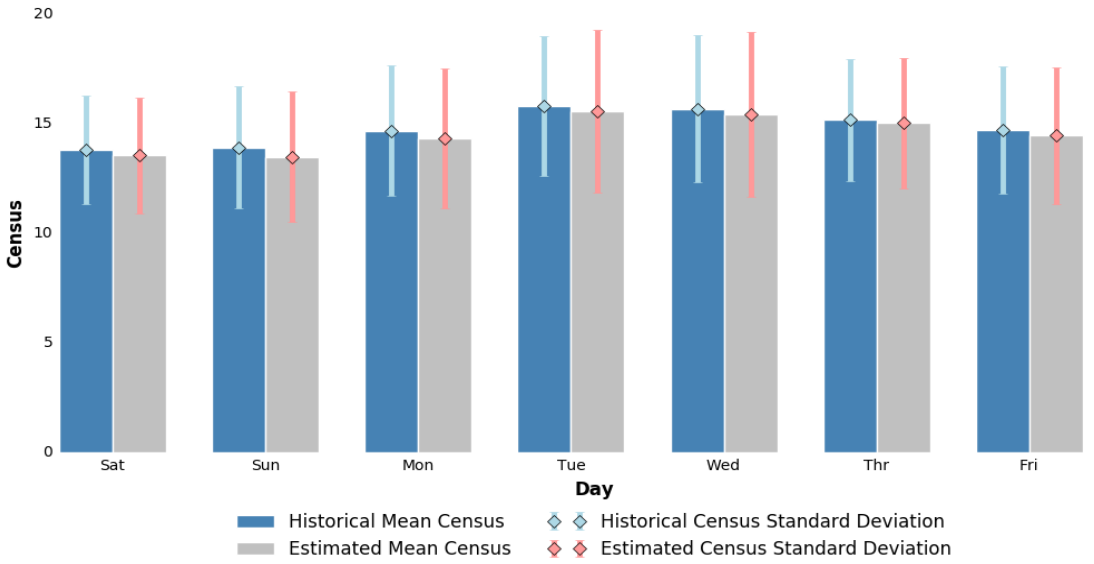

Figure 2.4: Daily historical and estimated mean and standard deviation of the census for one of the ICU units of hospital network (MICU).

depending on computationally expensive simulation models. We used the Gurobi optimization package through a Python interface to build and solve our MIP model. We applied our modeling methodology to patient data that spanned over 18 months to derive the optimal admission thresholds for different units under 97.5% service level. We consider two scenarios: 1) constant admission thresholds for different days of the week (OPT-C) and 2) dynamic admission thresholds varying by day of the week (OPT-D).

Since there is no historical record of hospital operational metrics under the optimal schedule, we compare the performance of our optimal admission policy both in constant and

Table 2.1: Accuracy of the weekly estimated mean census in comparison with observed data

	ICUs			IMCs					
	MICU	SICU	CCU	9E	6E	7E	6N	2E	GMB
Estimated μ	14.82	16.69	4.11	30.28	22.63	28.04	11.15	27.49	131.43
Observed μ	14.55	16.37	4.14	29.43	23.02	27.77	11.31	26.91	129.44
% Deviation	1.82	1.91	-0.73	2.80	-1.72	0.96	-1.43	2.10	3.42

Table 2.2: Simulation output for optimal admission policy vs. current admission policy (weekly)

	ICUs				IMCs					
Admissions	MICU	SICU	CCU	Total	9E	6E	7E	6N	2E	Total
Current Policy	14.23	15.31	6.87	36.41	32.12	31.26	37.71	7.05	22.07	130.21
OPT-C	16.05	17.37	7.14	40.56	34.02	33.77	39.31	8.51	24.44	140.05
OPT-D	19.52	22.19	8.11	49.82	35.93	35.54	42.15	8.91	25.32	145.85
Blockage										
Current Policy	1.261	1.513	0.464	3.238	1.082	1.195	1.188	0.725	0.911	5.101
OPT-C	0.702	0.717	0.304	1.723	0.831	0.922	0.926	0.519	0.628	3.826
OPT-D	0.692	0.731	0.273	1.696	0.845	0.897	0.934	0.487	0.595	3.758

dynamic settings with the hospital's current admission rule, using a high-fidelity simulation model. Our partner hospital currently follows an admission policy where patients with an MRM of 0.2 and higher are prioritized to be admitted to ICUs, and patients with an MRM higher than 0.07 are prioritized to be admitted to IMCs [45]. Although this is the guideline, in practice, some high risk patients are placed in GMB.

For the simulation modeling we follow a similar framework as discussed in [67]. However, we consider a more detailed unit transition probability metric for patients by conditioning on the arrival day and the arrival admission unit using $P^{u'}_{d'kui}(r)$ matrices as described in Section 2.3.2 and used in Theorems 2.1 and 2.2. Therefore, we can generate sample path of patients based on their arrival day and unit and build their stochastic trajectories until discharge. We use the historical data to estimate the transition probabilities as well as the distribution of arrivals, $\bar{F}_{d'ki}(a)$. We run the simulation for five years and report the results based on the last four years to account for a warm-up period of one year. We repeat the simulation 100 times and present the average results.

We validated our simulation model by considering the real-world hospital operations. We used almost two years of historical data to calibrate the model. To do so, we implemented the hospital's current admission rules in our simulation model and captured the key features

Table 2.3: Average length of stay of high risk patients for different admission policies (in days)

	OPT-D	OPT-C	Historical
Average Length of Stay	5.232	5.581	6.818

of the system such as the average unit's census by day of week, standard deviation of the unit's census by day of week, the average number of admissions for different units by day of week and the average length of stay per unit by patient class and mortality risk. We were able to see that the results of our simulation model nearly matches the actual hospital operations.

The results in Table 2.2 show that the optimal admission thresholds outperform the historical practice patterns. Under both constant and dynamic settings, our methodology improves the number of weekly admissions and weekly blockage for all ICU and IMC units. The OPT-C model increased the weekly average admissions to ICUs and IMCs by 11% and 7.5%, whereas the OPT-D model boosted that number by 37% and 12%, respectively. Also, the number of blocked patients in overall IMC and ICU units reduced by 33% and 34.5% in OPT-C and OPT-D models, respectively.

The improvement in admission and blockage shows that our methodology successfully assigned the patients to appropriate units upon patient arrival. The advantage of admitting higher risk patients to the higher levels of a hospital is twofold. First, on average it is expected to reduce the number of unplanned transfers. The reason is that in general beds (GMB in our model), higher risk patients are more likely to get sicker and require unplanned transfers to a higher level of care unit where they can be treated more effectively. Second, it should on average decrease the overall length of stay. It may be that higher level of care units, which offer a greater intensity of nurse care, are more effective at speeding up the healing of patients; however, it is beyond our scope to rigorously test that

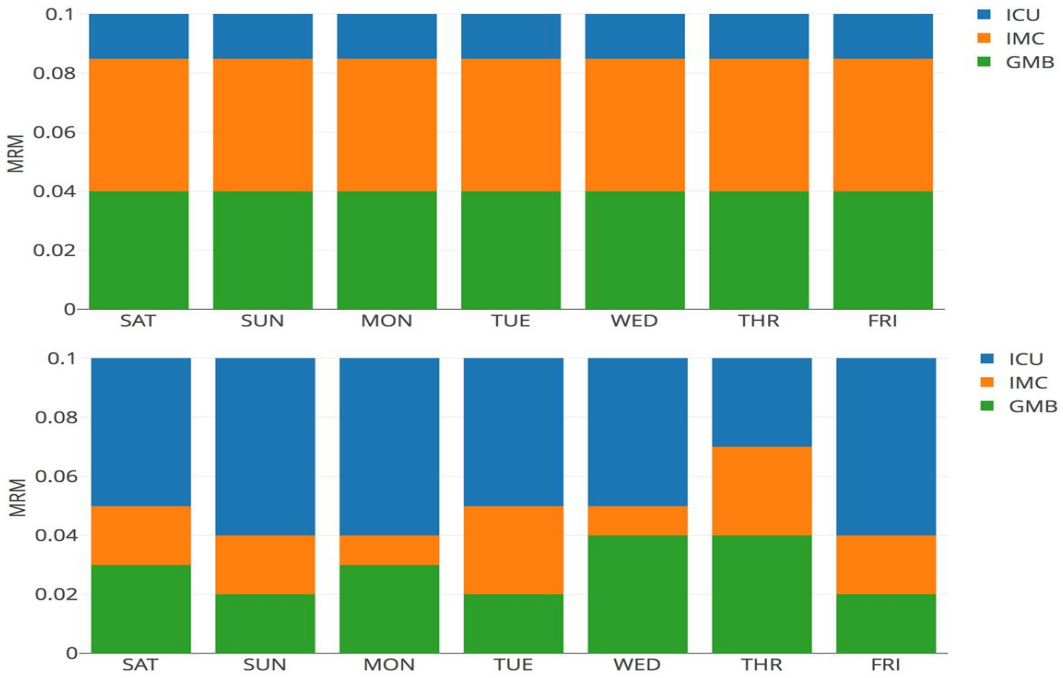

Figure 2.5: Top: Constant admission threshold based on mortality risk and day of week. Bottom: Dynamic admission threshold based on mortality risk and day of week.

hypothesis. Table 2.3 shows the difference in the average length of stay of high risk patients (MRM ≥ 0.06) for OPT-D, OPT-C and historical admission policies from the simulation results. The references cited in the Introduction, [13, 122, 123, 130], also align with both of the above relationships. In fact, incurring lower blocking rates and/ or having higher admission levels, something our results from optimization predicted, aligns with the above two relationships.

It is worth noting that the dynamic threshold policy enables the hospital to admit more patients to ICUs and IMCs than do static thresholds. By allowing dynamic thresholds, a hospital can maintain the desired service level and further increase the weekly admission volume to the higher levels of the hospital. This happens because the constant thresholds must be set to accommodate the busiest day of the week, whereas the dynamic thresholds can manage the tradeoffs and patterns observed in the stay of arrivals on various days to

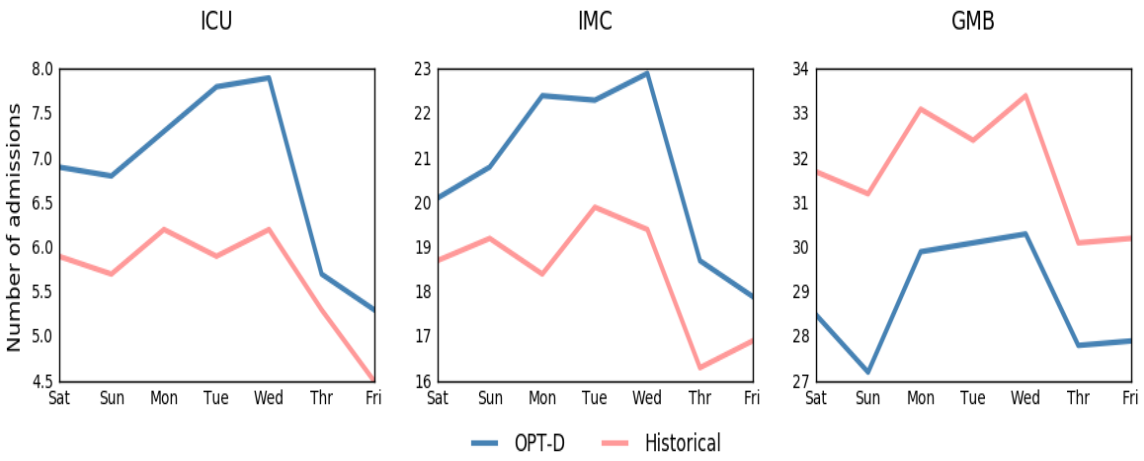

Figure 2.6: Daily average admissions to ICUs, IMCs and GMB for OPT-D vs. historical data. Mandatory placements are considered.

gain new performance. Please note that weekly patterns set by the surgical block schedule cause such patients to exhibit weekly patterns of need for beds and for higher level of care services. Surgical specialties vary in the length of patient stay and in patient need for higher level of care beds. Therefore, employing a dynamic admission policy for different days of the weekly plan is able to not only raise the long-run expected number of admitted critical patients but also to better exploit operational asymmetries across the days of the week.

To show how admission thresholds change in different models we set the daily risk threshold for different levels of care in a conservative fashion. For instance, to assign the admission threshold for ICUs on Saturday, we pick the highest threshold among all the ICU units and patient classes on that day. Figure 2.5 illustrates the significance of the improvements that result from the dynamic threshold model. OPT-D reduced the threshold for ICU admission on every day of the week. The lower admission threshold for both ICU and IMC indicates the higher availability of critical units for more high risk patients in the long run without compromising the service level.

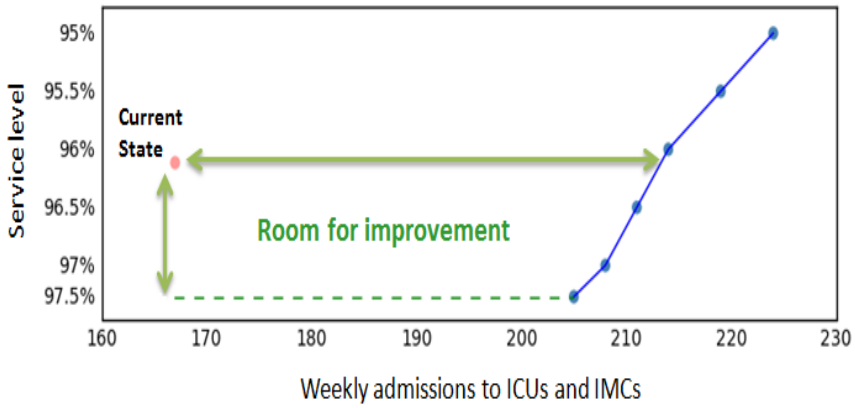

Figure 2.7: Pareto curve for service level vs. average weekly admissions to all ICUs and IMCs.

Figure 2.6 displays the average number of initial admissions by day of week for ICUs, IMCs and GMBs. Note that later transfers across units are not counted, although they are modeled. shows that OPT-D has selectively moved higher risk patients, who used to be admitted to general medical beds, to critical care units. Note that there is less of a difference on Tuesday, Wednesday, Thursday, and Friday. This is consistent with the fact that the units are the most full midweek. Our methodology empowers high level hospital strategies by providing data-driven operational prescriptions. Our approach can guide care providers to better understand the value of personalized admission control and how to make trade-offs.

Figure 2.7 shows the Pareto curve generated by solving the OPT-D model for service levels ranging from 95 to 97.5 percent (for the resources and arrival distribution of our partner hospital, higher service levels will result in infeasibility). This curve defines a meaningful boundary between utilization and accessibility that provides actionable insight for the hospital management system. Our optimization program delivered the operating curve automatically, with each point taking about 30 minutes to be generated. We should emphasize that our model can go far beyond what has been discussed in this section.

For instance, our model will be very insightful in capacity planning for critical care units where care providers can see how the system behavior changes by increasing or decreasing resources in critical care units. In general, our methodology can be seen as a predictive platform for a patient's future needs given her initial unit of admittance. Such a platform enables hospitals to test the robustness of their admission policies by easily generating and testing new arrival scenarios against their current admission rules.

2.5. Conclusion and Future Work

We introduced a novel perspective to address the admission policy for a network of critical care units. Although much work remains to be done to assess a patient's health risk, this paper takes advantage of one of the early measures of patient mortality risk to show how patient risk metrics can be used to rationalize and optimize the admission of patients given the limited capacity of higher level of care in a hospital. We developed a methodology that generates optimal admission policies for different ICUs and IMCs in a hospital network. Considering the formidable cost of critical care beds, our model can guide hospital admission policies towards a data-informed solution that guarantees a higher service level as well as higher accessibility.

Our suggested approach helps hospitals improve the effectiveness of their care delivery by ensuring that the neediest patients are admitted to the ICUs and IMCs at all times. By using a mixed integer program, we determined different thresholds for each unit and each day of the week. This lets us exploit the network effects inherent in the hospital and allow the policy to incorporate the weekly census pattern cycle and the stochastic care pathways revealed in the historical data.

It is essential to note that our proposed methodology is not limited to our partner hospital. It can be used by any hospital with a real-time health risk metric and a suitable

patient records database. We expect this approach to gain acceptance as hospitals deploy methods to assess patients' risks. Our method will diminish wait times to admit high-risk patients to higher level of care (ICU and IMC) by selectively assigning a more appropriate initial admission unit. Therefore, it has the potential to decrease mortalities caused by delayed access or non-access to critical care units. By incorporating stochastic location random fields, the method is based on an accurate prediction of patient unit/ward needs throughout her stay in the hospital.

Although our model carefully incorporates a congestion network model of hospital unit capacity, the capacity violation outcome can be different for hospitals with non-normal workload distribution. However, the normality of hospital workload has been a common observation in the literature. It is also worth mentioning that our methodology is most successful when the majority of ICU and IMC beds are not occupied by mandatory patients so we have more room for the optimal assignment of discretionary patients to the right units.

For future work, one direction is to consider dynamically updated patient mortality/health risk assessments, provided such longitudinal data are available. Dynamic risk measures can provide significant information regarding the patient recovery process that is valuable for improved admission decisions.

Chapter 3.

Optimization Under Uncertainty: Surgery Scheduling and the Promise of Individualization

3.1. Toward Individualized Care Delivery Planning: Big Data Analytics and Surgical Case Duration

3.1.1. Introduction

Surgeries are inherently stochastic processes, due to several medical and operational elements [77]. Operating rooms (ORs), however, are considered as the greatest source of revenue and cost for hospitals [48]. Therefore, efficient ORs are crucial in delivering hospitals' operational goals and better patient outcome [87, 94]. To achieve higher performance in surgical rooms and to mitigate waiting and idling times for both patients and staff, we require reliable estimation of case duration. This study offers a novel individualized

perspective to address case duration prediction problem by leveraging the entire patient-hospital data available prior to the surgery. We provide a clear roadmap to actually improve the care delivery system through the power of big data analytics [125].

OR utilization is a significant factor in hospital revenue management. Under-utilization of operating rooms costs $600 per hour and it is estimated that over-utilization is even more expensive [104]. Such operational flaws in ORs will enhance the pressure on the overall hospital network which imposes higher risk on patients and increases dissatisfaction among the staff. Therefore, hospitals constantly seek new opportunities to control OR costs and elevate the quality of surgical services by identifying significant sources of uncertainty in estimating the duration of surgery [129].

Traditionally hospitals adopted rudimentary point estimate approaches to predict the case duration in operating rooms. Sample mean and median are among the methods that are employed by a majority of hospitals. Such methods, due to a high level of aggregation and their natural simplistic assumptions, are prone to lack of fit and fail to provide reliable predictions. Introducing a time component, some hospitals stepped a little further by incorporating the median of the last n similar cases by the same or similar type of surgeons [91]. However, because of the complexity of defining similarity among different cases as well as the exclusion of key information, these practices are highly biased. Surgeon estimate is also another way by which hospitals try to control the variability of their OR schedules. Although surgeons' opinions convey valuable information which is a combination of experience and subject matter expertise, this approach also lacks a comprehensive view of all aspects of the problem and are shown to be less accurate than statistical methods [83]. Hybrid methods consisting of both surgeon and statistical point estimates are also studied in various hospital settings, but negligible improvement was observed [143].

To address these issues and to provide a framework that will not only generate reliable

predictions but also deliver insightful statistical decision support, we propose a comprehensive data-driven methodology. This study, in terms of patient-hospital data, offers the most inclusive and elaborate approach for case duration prediction to date. We include various layers of patient and hospital interactions as well as numerous operational measures to fully capture the system dynamics. To be able to define and harness the variabilities in case duration, we focus on four main categories: 1. patient related variables, 2. hospital and staff related variables, 3. operational variables and 4. temporal variables. Considering several key contributors, we will exploit all these categories to the full extent to engineer powerful and informative features. We then use these features to build several statistical models which not only yield state-of-the-art prediction accuracy but also deliver empowering data-informed toolbox to guide hospital's operational decision making in surgery sequencing and scheduling. We compare the results of these models and provide actionable managerial insights.

Surgeries are consisted of several related activities that complement one another. In general, surgeries are defined by one or multiple current procedural terminology (CPT) codes. It is well-known that CPT codes convey useful information about the overall surgery operations and are, therefore, helpful in estimating case duration [50]. However, CPT codes have grown immensely large in size and can easily comprise a cohort of hundreds of unique codes for each service. This issue complicates the use of CPT codes as they can generate significantly sparse data structures [85]. In order to take advantage of information stored in CPT codes without incurring unnecessary model complexity, we input Resource-Based Relative Value Scale (RBRVS). RBRVS is a mechanism by which physician fees are determined [70]. These values are intended to measure the amount of physician work for each of the procedures. Therefore, they are finely linked to CPT codes. This study, to the best of our knowledge, is the first work to investigate the importance of RBRVS in

predicting case duration.

We also do not restrain our predictive power to the available structured data rather we embrace and digest all the unstructured raw inputs that are collected prior to the surgery. From several discussions with surgeons, anesthesiologists and staff, we realized that there is invaluable amount of information in documented patients' text records that can fairly reveal the complexity of a surgery. Therefore, to better differentiate surgeries and to further understand the influential factors in defining case length variations on an individualized case-by-case fashion, we explore multiple text notes from surgeons and radiologists. We introduce generalizable methods for feature extraction from large textual data that can be easily applied on various services within a hospital. We then incorporate the output of the text analysis in our predictive model to take full advantage of the available raw data in a conveniently implementable framework.

With the increasing abundance and accessibility of Electronic Health Records (EHR), this study casts light on new directions for the inevitable application of big data to healthcare [4, 101]. This paper contributes a personalized and generalizable statistical modeling methodology that is capable of offering much more than point estimates of surgery duration. Unlike the previously studied approaches, we will investigate the uncertainty level of our predictions and show how valid prediction intervals can be utilized to inform better sequencing of daily surgeries. We will validate our results via extensive experiments on rich data sets of Memorial Sloan Kettering Cancer Center (MSKCC).

3.1.2. Literature Review

During past few years, researchers have used tools from statistics and mathematics to build predictive models for surgical case duration [52, 53, 102]. However, the scope of these studies is very much limited to a few quantitative and categorical predictors [129].

Eijkemans et al., for instance, included team, patient and session characteristics in their modeling perspective and they noted that patient characteristics are essential predictors in the model [56]. Stepaniak et al. also investigated the dependence of procedure times on surgeon and team characteristics and reported team composition, experience and daytime as the most often significant factors [127]. Many also investigated the effect of procedures and anesthesia and showed the presence of such parameters in the model is statistically significant [65, 148]. For example, Strum et al. tested goodness of fit for a normal and log-normal distribution on surgical procedure times based on CPT codes and anesthesia type [128].

Although common statistical methods are helpful in developing a solid foundation to identify key variabilities, they have been outperformed by novel machine learning (ML) tools in terms of accuracy. Recently, several researchers have embarked on employing ML algorithms in healthcare operations [5, 88, 108]. Although recent studies shed light on the impact of preoperative data in defining surgical uncertainties, ML techniques have not been utilized extensively in case duration prediction area [24]. Also, those who investigated this topic restrained the scope of their studies to only a few categories of variables [43, 46]. For instance, Li et al. only focused on CPT codes and ignored all patient, surgery team, temporal and operational variables [85]. In a more comprehensive attempt, Kayics et al. studied operational, temporal and behavioral factors; however, they completely disregard patient-specific data which is a key contributor to case duration variations [76]. In this paper we go beyond what has been offered in the literature by including all accessible data prior to the surgery.

As the availability of unstructured data is quickly growing along with the structured data, more researchers are moving towards text mining information retrieval algorithms in healthcare related topics [36, 66, 74]. Despite the abundance of readily obtainable textual

data in hospitals' databases such as surgeon's initial consults and pre-surgical notes as well as radiologist's recommendations and reports, there has been no conclusive study to explore such resources in the realm of case duration analysis. In this study, as the first attempt to incorporate the power of text analytics in estimating the duration of surgeries, we show that generalizable text mining techniques can be applied to automatically extract informative features from various textual data repositories.

It is well-known that OR management heavily depends on estimating surgery duration [56]. Therefore, case duration prediction has also been studied extensively from the perspective of operating room planning and scheduling [48, 49, 92]. However, these studies mainly focused on point estimates, moments and discrete set of scenarios for case duration and did not provide prediction intervals which are crucial in operational decision making process. There are very few examples in the literature focusing on the uncertainty bounds of predictions. Dexter et al., for instance, suggested a Bayesian method to generate prediction limits for surgery duration [51]. But, their approach is parametric and involves certain assumptions on the unknown distribution of case duration. In contrast, we propose a fully data-driven methodology to generate prediction intervals for individualized case duration that can be further utilized as a decision support tool in an OR planning environment.

3.1.3. Data

In this study, for initial exploratory analysis, we consider the surgeries across all services performed at MSKCC within the last six years (2010 - 2016). To describe our statistical model, then, we focus on Gynecology (GYN) service. However, we further show that our methodology can be easily generalized and applied on all services throughout the hospital.

MSKCC is one of the leading cancer centers in the world which is constantly improving cancer treatment process by applying state-of-the-art scientific methodologies in both

Table 3.1: Operational data for different services in the hospital

Service	Surgeries	Surgeons	Inpatient (%)	Case Duration (SD)
BRE	21,424	15	3,397 (15.9)	118.8 (97)
CRS	7,437	7	6,560 (88.2)	239.2 (160)
GMT	8,518	11	4,914 (57.7)	157.0 (109.1)
GYN	12,281	15	5,699 (46.4)	201.4 (123.4)
HEP	6,605	7	6,190 (93.7)	226.0 (104.1)
HNS	11,529	13	6,617 (57.4)	197.9 (149.4)
NEURO	5,638	11	5,614 (99.6)	256.3 (115.4)
OPTH	1,700	3	554 (32.6)	79.9 (41.7)
ORTHO	5,213	6	4,030 (77.3)	240.4 (177.3)
PEDS	3,100	3	1,895 (61.1)	195.5 (161)
PLA	9,665	8	1,768 (18.3)	151.9 (110.4)
THO	13,456	14	9,442 (70.2)	170.9 (119.9)
URO	23,176	49	9,428 (40.7)	153.5 (129.2)

medical and operational aspects of its care delivery system. MSKCC surgical operations take place across thirteen different services and in a total of 40 different operating rooms. We work with a total of 129,742 distinct surgeries in all services throughout the study horizon that include all inpatient and outpatient surgeries. These surgeries range from one-procedure to twenty-four-procedure cases with 93% of them containing fewer than 6 procedures. Based on CPT codes, an overall of 2,346 unique procedures are observed and 1,730 distinct primary procedures (the procedure with the highest RBRVS in a case) are reported. Table 3.9 provides a detailed summary of hospital operations across all services.

Due to the large number of unique procedures, the CPT codes do not provide a practical set of predictors as they can considerably boost the sparsity of our dataset. Therefore, we use RBRVS measures which are continuous numerical values that are meant to estimate the effort required for performing a particular procedure. For each case, we consider a

Table 3.2: Demographics of the patient data

			Race, Number of Observations (%)			
Service	Age (SD)	Female (%)	White	Black	Asian	Other
BRE	55.6 (12.6)	21244 (99.2)	16,596 (77.5)	1,953 (9.1)	1,496 (7.0)	1,379 (6.4)
CRS	59 (13.8)	3405 (45.8)	6,073 (81.7)	491 (6.6)	444 (6.0)	429 (5.8)
GMT	60 (14.9)	3,829 (45.0)	7,036 (82.6)	503 (5.9)	481 (5.6)	498 (5.8)
GYN	55.2 (13.6)	12,279 (100)	9,927 (80.8)	824 (6.7)	860 (7.0)	670 (5.5)
HEP	61.4 (13.2)	3155 (47.8)	5,545 (84.0)	314 (4.8)	398 (6.0)	348 (5.3)
HNS	55.3 (16.2)	6,065 (52.6)	9,346 (81.1)	480 (4.2)	875 (7.6)	828 (7.2)
NEURO	55.1 (17.5)	2,737 (48.5)	4,659 (82.6)	345 (6.1)	357 (6.3)	277 (4.9)
OPTH	54.6 (23.5)	827 (48.6)	1,501 (88.3)	69 (4.1)	54 (3.2)	76 (4.5)
ORTHO	46.9 (21.3)	2,783 (53.4)	4,247 (81.5)	397 (7.6)	306 (5.9)	263 (5)
PEDS	11.6 (8.3)	1,375 (44.4)	2,304 (74.3)	407 (13.1)	239 (7.7)	150 (4.8)
PLA	51.4 (11.3)	9,124 (94.4)	7,898 (81.7)	839 (8.7)	464 (4.8)	464 (4.8)
THO	63 (13.2)	6,646 (49.4)	11,409 (84.8)	670 (5.0)	732 (5.4)	645 (4.8)
URO	63.1 (13.6)	5196 (22.4)	19,544 (84.3)	1,466 (6.3)	799 (3.4)	1,367 (5.9)

combination of minimum, maximum, mean and sum of RBRVS measures as well as the total number of procedures to best represent the information stored in CPT codes without actually using them.

Due to the diverse nature of MSKCC outreach, the patient population encompasses a mixed distribution of ethnic and socioeconomic groups. Table 3.10 summarizes key demographics of patient data for various services of the hospital.

The focus of this study is estimation of the overall time a patient spends in the OR, that is the time interval between patient's move-in and move-out (toes-in-toes-out). Since the set-up and clean-up times are normally considered as standardized tasks, toes-in-toes-out constructs the most variable time window from the OR utilization perspective. MSKCC is already using a filtering system to estimate the length of surgeries. Currently, this is done by leveraging a feature within MSKCCs implementation of Epic's OpTime module, which takes the median of a set of historical cases that are matched by CPT codes, surgeon, and location. If there is no past observations, they select from similar surgeries performed

by other surgeons. Although the current method seems to be inclusive of several layers of information, it clearly lacks a holistic view of system dynamics on a patient-by-patient basis. Because the functions available in OpTime require a relatively rigid set of criteria for matching, along with the variability inherent in surgical procedures, over 50% of surgeries have an actual duration outside of ± 30 minutes from the original estimate. To overcome the limitations of the current approach and to deliver a personalized framework, we employ structured and unstructured data to input every piece of information available prior to the surgery without any aggregation.

The available structured data for this study can be classified into main groups of patient, surgery team, equipment, operational and temporal data. On the patient side, we include age, gender, race, BMI, treatment history and Elixhauser comorbidity measures. On the other hand, to expand our information acquisition outreach, we incorporate data on case RBRVS measures, equipment, surgeons, operational specifications of surgeries, etc. A detailed list of structured data used in the final predictive model is given in Section 3.1.4.

Since most of the patient individual characteristics are recorded in unstructured format, focusing only on structured data can considerably undermine the personalization. Unstructured data refers to the raw records available from patient and hospital that does not have a pre-defined data model in a relational database. There are different instances of unstructured data in healthcare such as images, videos and texts. Due to the irregularity and inconsistency in their structure, these types of data are usually ignored in statistical modeling. However, after several discussions with surgeons, anesthesiologists and clinicians it became clear that significant amount of information is stored in text documents. Such information conveys certain specifications, individual characteristics and particular treatment requirements of each patient. In this study, therefore, we focus on the textual data recorded and maintained in various notes as in initial consult, follow-up updates,

radiologist report and generally all textual records available on patient profile prior to the surgery.

Table 3.3: Features extracted from structured data

Category	Features
Hospital	The primary surgeon, experience of the primary surgeon in years, number of publications of the primary surgeon, surgery room, number of panels in the surgical case, the sum of all RBRVS measures for the case, the mean of all RBRVS measures for the case, the maximum of all RBRVS measures for the case, the minimum of all RBRVS measures for the case, number of procedures for the surgical case, number of robotic procedures, the mean duration of the last 5 similar cases based on the primary procedure, the number of times the surgeon had a similar surgery in the 30 days, the number of times the surgeon had a surgery in the 30 days, the number of times the surgeon had a similar surgery in the past, the number of times the surgeon a surgery in the past, whether or not a particular equipment was needed for the surgery, the number of required equipment
Patient	Elixhauser comorbidity measures such as obesity, depression, etc., Body Mass Index (BMI), weight, race, gender, age, whether the patient is inpatient or outpatient, the number of days spent in the hospital prior to the surgery, the number of times patient underwent chemotherapy, the number of times patient underwent radiation therapy, the number of times patient underwent surgeries
Operational	The number of days between the day of surgery and the day the surgery was scheduled, the number of cases assigned to the surgeon on the day of surgery, the number of cases assigned to the surgery room on the day of surgery, the number of cases scheduled on the day of surgery, the sequence number of the surgery on the surgeon's schedule, the sequence number of the surgery in the surgery room
Temporal	The scheduled time of the surgery, The weekday of the surgery, the month of the surgery, the year of the surgery

3.1.4. Methodology

Our purpose in this research is not only to improve the accuracy of the current predictions but also to develop insightful statistics associated with each of the estimates. We are interested in capturing the contribution of various factors in defining the variations of surgery duration and, therefore, we consider a multivariate approach. Considering that the target dependent variable is continuous, we use regression techniques to estimate the surgical case duration. In this section we describe the crucial steps in building our statistical model from data pre-processing to construction of the predictive models.

Feature Engineering

Data preparation is an essential phase in any machine learning task. As described in Section 3.2.5, we receive data in various forms and types and Without a systematic approach to generate informative features to avoid noisy, irrelevant or redundant data, the process of knowledge discovery becomes very complicated and hard to generalize. Therefore, we have to develop an automated mechanism that transforms the raw data into meaningful predictors for our models.

The raw data table available for this study consists of different interactions of the patients and the hospital with all required time stamps, certain test results, static information of the staff and the hospital facilities and detailed specifications of each surgical case. In the pre-processing phase, we used the raw data to generate a multitude of new features in different groups. Table 3.11 provides a complete list of features derived from the structured data.

Feature Extraction and Dimensionality Reduction for Textual Data

After consulting the surgical staff, we realized that patients' clinical notes may contain certain words, medical terminologies, treatment options, etc. that indicate complexities and uncertainties of a particular surgery. In fact, it is very common that the existence of specific terms on the patient profile differentiate his/her case from other similar surgeries. For instance, a surgical case in which the patient has an implanted "surgical mesh" tends to be more challenging and time consuming. Based on this idea and given the fact that we are not concerned with ordering and semantics of the words, we only focus on the frequency of terms in the textual records of patients.

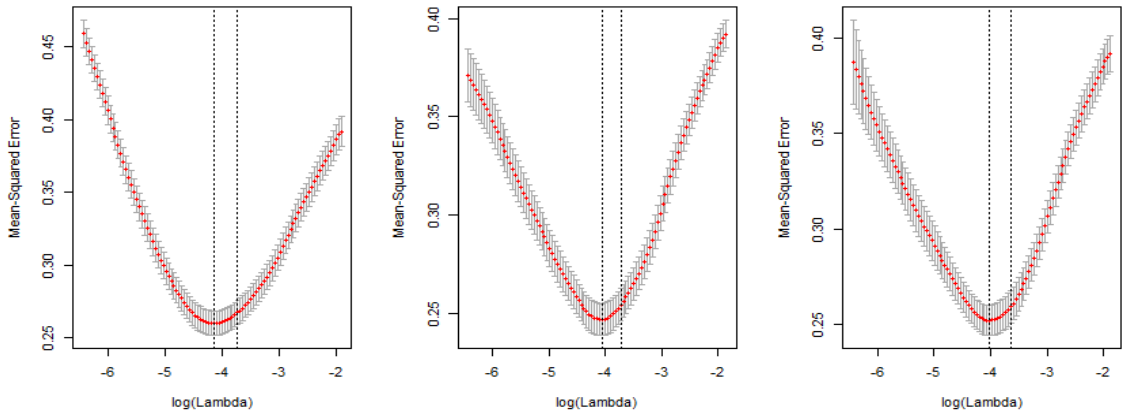

Figure 3.1: Mean squared error for different λ values. From left to right: uni-grams, uni-grams and bi-grams and uni-grams, bi-grams and tri-grams

In order to input the clinical notes to our modeling framework, for each patient we concatenate all pre-surgery records and build one document from initial consults, follow-up notes, etc. The final output of this process is a corpus of textual documents. We then incorporate the bag-of-words (BOW) model to transform these documents into a structured format. After initial data cleaning, all distinct terms of the corpus are identified and BOW transforms each document by constructing a term frequency vector. The terms are defined

as different contiguous sequence of words in the corpus. In our case, we focus on uni-grams, bi-grams and tri-grams which refer to sequences of one, two and three adjacent terms, respectively.

However, this method results in a very large feature space. We, therefore, further process the output by removing stop words, non-alphanumeric characters, terms that are seen in a majority of documents and terms that are so rare that no statistical evidence can be inferred from them. We also reassign new weights to each term based on the frequency of terms in the document and the overall corpus. This approach is called term frequency-inverse document frequency (tf-idf) and is a standard practice to adjust for the fact that some terms are generally more frequent.

With the pre-processing tasks described above, we end up with a few thousands of term

Figure 3.2: Wordcloud of the significant uni-grams in GYN corpus derived from the bootstrapped Lasso model

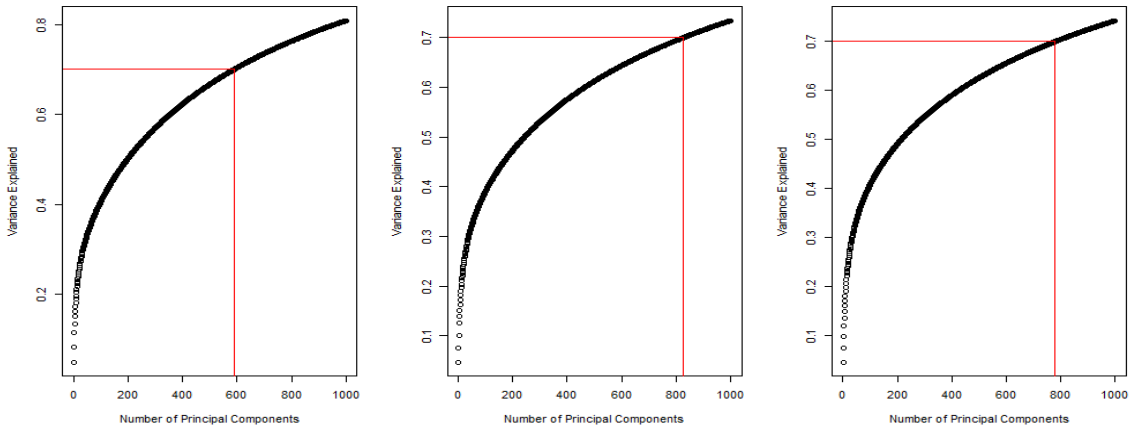

Figure 3.3: Explained variance of the corpus data for different number of principal components. From left to right: uni-grams, uni-grams and bi-grams and uni-grams, bi-grams and tri-grams

features for any of the services in our hospital. Most of these terms are still statistically insignificant in explaining the variations of surgical case duration. Also, using such a high dimensional data structure to complex machine learning algorithms is computationally costly. Therefore, in order to automatically identify the most important terms or combination of terms in our corpus and to shrink our feature space, we use two different strategies.

In the first approach we implement Lasso Regression (LR) as a feature selection tool. Lasso Regression is an interpretable statistical model that minimizes the residual sum of squares while forcing the sum of the absolute value of the dependent variables' coefficients being less than a constant [131]. This algorithm tends to construct simpler models by dropping insignificant predictors which is, given the magnitude of our feature space, a desirable quality for our application. For each service, we regress the surgical case duration variable on the the feature space of distinct terms of their associated corpus. The Lasso then performs L_1 shrinkage which generates sparse models by forcing the coefficients of insignificant features to be zero [131].

To generate a stable list of significant textual terms we run the Lasso model on 100 bootstrapped samples of our dataset. The sparsity of the final model in each bootstrap iteration is controlled by a penalty factor that is assigned to the variable coefficients which is tuned by applying a 10-fold cross validation. Figure 3.1, for instance, illustrates how the mean squared error (MSE) on one bootstrap validation set changes with respect to the penalty factor (λ) for the Lasso model on GYN corpus. To account for estimation error in model selection, it is recommended to use the λ that generates an MSE within one standard deviation ($\lambda.1se$) of the minimum MSE ($\lambda.min$) [60]. Accordingly, we insert tf-idf measures of the terms with non-zero coefficients in $\lambda.1se$ model that have non-zero coefficients in more than 50% of the bootstrap replicates into our final statistical learning model. Figure 3.2 illustrates the significant terms in GYN corpus where the font size is proportional to the number of times each term ended up in the final model on the bootstrapped samples.

Our second strategy is to apply Principal Component Analysis (PCA) to reduce the size of our feature space while maintaining a certain level of variability in our corpus. Figure 3.3 shows how well the compressed data represents the original corpus as we add more principal components. To ensure a meaningful dimensionality reduction, we focus on the leading principal components that capture 70% of the textual data variations. Table 3.4 depicts different phases of dimensionality reduction (DR) for the textual features in GYN service with bootstrapped Lasso and PCA. Using either of these methodologies, one can programmatically output a significantly smaller feature space from unstructured medical notes which can be further combined with previously discussed structured data. Both mechanisms are also easily generalizable to all services across the hospital.

Table 3.4: Textual feature extraction for GYN corpus (uni-grams, bi-grams and tri-grams)

DR step	Uni-grams	Uni & bi-grams	Uni, bi & tri-grams
Raw data	24265	392528	1177118
Pre-processed data	6430	30720	50505
Bootstrapped Lasso ($\lambda.min$)	347	408	383
Bootstrapped Lasso ($\lambda.1se$)	221	311	262
Principal Components	590	826	779

Statistical Learning Models

After data preparation in both structured and unstructured formats, we can move forward with an in-depth analysis of statistical models for the surgical case duration problem. As for the main predictive model, we chose to use two different machine learning algorithms: 1. Gradient Boosting Machine and 2. Random Forests.

Gradient Boosting Machine (GBM), enjoys from an iteratively additive structure where in each iteration a weak learner is introduced to compensate the shortcomings of the existing weak learners from previuos iterations [59]. This algorithm is very powerful for regression purposes as it fits an additive model in a forward stage-wise fashion where shallow decision trees are built in sequential order complementing one another. But, GBM is prone to over-fitting and has several hyper-parameters that complicate the tuning process of this algorithm.

Random Forests (RF), on the other hand, is another widely applied ensemble method that uses decision trees as base learners [26]. In comparison with GBM, RF employs fully grown decision trees that are generated in parallel rather than sequentially. Additionally, RF is much easier to tune and, due to its random behavior, is more robust to over-fitting and superior in generalization performance.

Although ensemble methods have shown promising results for high-dimensional regres-

sion problems and are commonly used in predictive modeling, there is little known about the statistics of their predictions. To have a concrete operational view, however, it is critical to understand the prediction intervals and uncertainty bounds associated with final estimates.

Quantile Regression

Given the covariates extracted from the data, ensemble algorithms make accurate approximation of the conditional mean of surgical case duration. However, in addition to the conditional mean, we are interested in learning about the full conditional distribution of surgery length. In fact, for operational planning, it is vital to know the uncertainty bounds around predictions. For instance, it might be possible to estimate the duration of a particular surgery to a higher accuracy than another surgery on a given day. This information delivers a decision support tool to improve case sequencing and scheduling on a daily basis.

Quantile regression is able to provide a more comprehensive view of the response variable (case duration) [80]. The conditional distribution function of case duration is defined as the probability that given a feature set $X = x$, the duration variable D is smaller than $d \in R$,

$$F(d|X = x) = P(D \leq d|X = x). \tag{3.1}$$

Then, given the feature set $X = x$, α-quantile $Q_\alpha(x)$ will be constructed such that the probability of D being less than $Q_\alpha(x)$ is equal to α,

$$Q_\alpha(x) = inf\{d : F(d|X = x) \geq \alpha\}. \tag{3.2}$$

Quantiles, therefore, generate more information about the fluctuations and spread of surgery length for every individual. We can use quantile regression to construct prediction

intervals. For example, if we do not want to over-estimate case duration with high probability, we can easily build a 95% prediction interval for any future surgical case given its associated feature set X by

$$PI(x) = [Q_{0.025}(x), Q_{0.975}(x)]. \tag{3.3}$$

In order to generate accurate and fully non-parametric prediction intervals while incorporating more complex learning models such as RF, we use Quantile Regression Forests (QRF) [96]. QRF follows quite a similar algorithm as Random Forests. But, instead of approximating the conditional mean, it uses the weighted observations of response variable to approximate the full conditional distribution.

3.1.5. Validation and Numerical Results

To implement and test the models for surgical case duration prediction, we split the data in GYN service into two parts. The chronologically first 80% of surgical cases comprise the training set used to calculate and tune the required parameters for each model. The remaining 20% of the surgical cases will be maintained in the test set which will be used to test the performance of the developed models. We build and compare GBM and RF models with and without textual data with different combinations of PCA and bootstrapped Lasso on textual features. In order to train the GBM and RF models and to tune their hyper-parameters we implement a 10-fold cross validation. Table 3.5 shows the accuracy of the current method in comparison with our individualized models. We report the Mean Error (ME), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE) to evaluate the performance of these models on the test set. Based on the RMSE measures, the RF model that is trained on the bootstrapped

Lasso textual features of uni, bi and tri-grams has the highest accuracy. The improvement in RMSE is more than 23%. The RF Tri-grams model generates an MAE of 44.42 minutes which indicates that our suggested model reduces the prediction error of each surgical case duration in GYN service by almost 12 minutes on average. We should also note that the inclusion of textual data in RF contributes to a reduction of more than 2 minutes in average estimation error.

The impressive performance of our modeling framework in achieving greater accuracy is partly due to mitigating extreme errors of underestimation. The current method is based on the median duration of previous similar operations and it tends to predict shorter duration for majority of cases, see Table 3.6. This mainly happens because the distribution of the surgical case duration is right-skewed and, therefore, predictions that are generated by minimizing deviations from the median are smaller than predictions that are generated by minimizing deviations from the mean.

In order to gain a better understanding of how our individualized setting improves the quality of predictions, we have to explore the impactful variables of the final RF tri-grams model. RF rates the importance of variables through a permutation test on different variables in the out-of-bag (oob) samples. The idea is that permutation on important variables will degrade the accuracy of prediction more significantly than permutation on non-important variables. Table 3.7 summarizes the significant features used by the random forest algorithm in different categories. The features are ordered by their importance in each category. One interesting observation regarding the selected textual features is that they convey very important information that could have been lost, if we only considered the structured data. For instance, a surgical case booked with HIPEC will likely require an extended open procedure with extensive intra-abdominal resections followed by heated intra-peritoneal chemotherapy and sometimes a reconstructive procedure of bowel followed

Table 3.5: Model performance comparison for GYN service

Models	ME	RMSE	MAE	MAPE
Current Method	-33.71	83.64	56.33	43.11
RF No Text	18.19	66.52	46.78	22.07
GBM No Text	12.16	65.72	45.01	22.14
RF Uni-grams	14.60	65.45	45.23	21.63
GBM Uni-grams	12.88	67.90	45.30	21.67
RF Bi-grams	15.84	66.41	45.80	21.76
GBM Bi-grams	13.33	72.56	47.73	22.75
RF Tri-grams*	**13.04**	**64.29**	**44.42**	**21.66**
GBM Tri-grams	13.38	69.82	46.27	22.09
RF Uni-PCA	14.73	67.22	46.26	22.12
GBM Uni-PCA	15.96	65.04	45.06	21.52
RF Bi-PCA	14.77	67.78	47.00	22.49
GBM Bi-PCA	14.61	66.32	46.30	22.66
RF Tri-PCA	15.23	67.38	47.15	22.52
GBM Tri-PCA	15.11	66.94	45.90	22.03

by closure. Cases who have undergone neoadjuvant chemotherapy and now brought for surgery are usually well selected cases that have responded to chemotherapy and the burden of disease is reduced by the chemotherapy and will likely have a more abbreviated operation than upfornt primary debulkign surgery.

The required data for developing individualized models are accessible well ahead of the surgery in various services across the hospital. This enables our methodology to be easily generalizable to other services. In order to test the universal performance of our

Table 3.6: Error quantiles for the best model in GYN service

Models	10%	25%	50%	75%	90%
Current Method	-117.00	-64.50	-26.00	3.00	35.00
RF Tri-grams	-44.91	-5.15	17.69	42.85	72.27

approach, we compare the accuracy of the best model in each service with the current estimation method on the test set. To find the best model in each service we follow the same steps as we described for GYN. Table 3.8 summarizes the performance results of the

Table 3.7: Significant features in each category (GYN Service)

Category	Significant Features (ordered by importance)
Hospital	The mean duration of the last 5 similar cases based on the primary procedure, number of panels in the surgical case, the sum of all RBRVS measures for the case, the maximum of all RBRVS measures for the case, number of robotic procedures, the number of required equipment, the number of procedures for the surgical case, Required equipment "ESU Force FX", Required equipment "ESU LIGASURE", Required equipment "PLASMAJET", Required equipment "XI Robot", Required equipment "ESU TISSUELI", the number of times the surgeon had a similar surgery in the past
Patient	Whether the patient is inpatient or outpatient, Body Mass Index (BMI), weight, age
Operational	The number of cases assigned to the surgery room on the day of surgery, the number of cases assigned to the surgeon on the day of surgery, the sequence number of the surgery in the surgery room, the sequence number of the surgery on the surgeon's schedule
Temporal	The scheduled time of the surgery, the year of the surgery
Textual	"debulking", "advanced-ovarian-cancer", "pelvic-mass", "omental", "omental-caking", "ovarian-cancer", "peritoneal-carcinomatosis", "supracervical", "exenteration", "cervical-cancer", "perihepatic", "robotic", "carcinomatosis", "cytoreductive", "retroperitoneal", "hydronephrosis", "hipec", "solid-cystic", "radical", "obese", "retroperitoneal-lymph", "effusion", "myomectomy", "fibroid", "hemipelvis", "hyperlipidemia", "neoadjuvant-chemotherapy-plan", "unresectable", "carboplatin", "metastatic-ovarian", "spread", "serous"

Table 3.8: Universal performance report; H.: Hospital current method, M.: Individualized model

Services	H.RMSE	M.RMSE	RED.(%)	H.MAE	M.MAE	RED.(Mins)
BRE	38.20	33.01	13.59	26.23	19.76	6.47
CRS	118.15	86.41	26.86	83.13	57.12	26.01
GMT	83.18	66.41	20.16	55.95	45.10	10.85
GYN	83.64	64.29	23.13	56.33	44.42	11.91
HEP	98.30	61.69	37.24	73.00	43.97	29.03
HNS	85.39	62.42	26.90	61.7	42.11	19.59
NEURO	117.48	86.54	26.34	79.27	54.59	24.68
OPTH	35.60	28.53	19.86	26.31	21.85	4.46
ORTHO	169.34	111.63	34.08	118.80	71.05	47.75
PEDS	106.56	70.13	34.19	71.28	47.22	24.06
PLA	64.77	48.00	25.89	45.90	32.64	13.26
THO	93.46	57.85	38.10	68.62	39.39	29.23
URO	56.88	39.26	30.98	38.20	25.67	12.53

best individualized model in each of the services available at MSKCC.

In addition to its higher accuracy, our methodology is capable of generating data-driven uncertainty bounds for each prediction. To collect quantile information of predicted values, we implement the quantile regression version of random forest (QRF) which for every leaf of every tree takes note of all observations and not just their average. Figure B.12 depicts the 95% prediction interval for a sample of 100 observations from the test set. The graph shows that fewer than 5 out of 100 actual durations fall out of the 95% prediction interval which supports the statistical performance of QRF [96].

Our modeling framework quantifies the bounds on stochasticity of surgical case duration in an individualized setting in which a comprehensive view of both patient and system

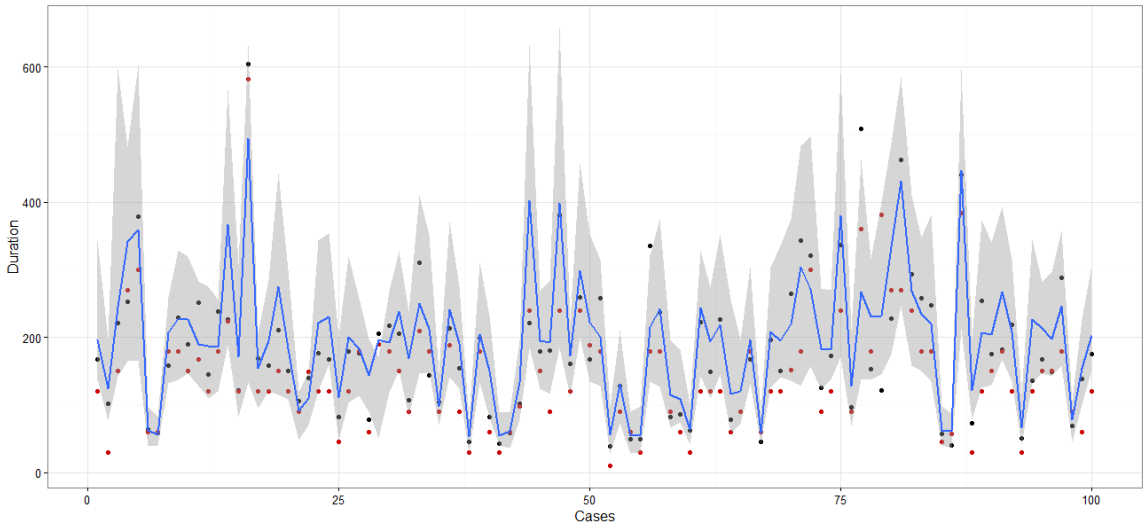

Figure 3.4: Prediction interval for a sample of 100 cases in GYN service; Blue line: predicted duration using RF model, Black dots: actual duration, Red dots: current planned duration, Gray margin: 95% prediction interval for RF model

is considered. Reliable information on the uncertainty of predicted durations provides a valuable data-informed toolbox for schedulers and operation managers. For instance, it is known that sequencing daily surgeries in order of increasing variance of durations reduces the schedule tardiness (overtime) [49]. Accordingly, using the results of our suggested model, one can easily rank surgeries based on their prediction intervals.

It is important to note that the QRF model is able to provide comprehensive information regarding the distribution of each individual surgery which is well beyond the mere uncertainty bounds discussed above. In the next section of this chapter we propose analytical methodologies to incorporate such individualized distributional predictions into the surgery scheduling problem under uncertainty.

3.2. An Individualized Learning Methodology for the Surgery Appointment Scheduling Problem

3.2.1. Introduction

Uncertainty is an inherent component of hospital operations. Surgery, in particular, suffers from a high rate of variability due to many factors such as co-morbidities, unforeseen complications, the patient's unique anatomy and several other medical and operational elements [77]. Consequently, operating rooms (ORs) are often affected by significant schedule volatility and poor utilization [48, 94, 104]. Ambivalently, ORs are well known to be both the greatest sources of revenue and cost for hospitals [124].

To improve the management of ORs, and in particular the setting of case start times, one of the crucial modeling components is the surgical case duration [56]. Its variability crucially affects the quality of decisions and require proper consideration. For instance, surgery durations that are short on average but wildly vary may harm a decision more than durations that are long but more stable [48]. This motivates researchers to study surgery scheduling optimization where the surgery durations are stochastic [48, 49, 92]. These studies, however, consider the duration distributions on a population-level (i.e., each patient is treated equally as coming from a single "generating" pool or from a number of different categories/types), mostly because breaking down patient classes would reduce substantially the effective data size and lead to poor fitting [48, 90].

Given the limitations of these approaches and advances in data analytics, several researchers have instead studied machine learning methods to predict surgical cases using patient-specific data [76, 88, 117]. While these methods are shown to be effective in estimating the surgical case duration, they, like most machine learning methods, are designed as point estimators, e.g., mean response. Therefore, they are unable to provide significant

information on the distribution of the case duration. However, when it comes to surgical scheduling, distributional variability plays a major role. As a result, these studied methods can fall short of making the best decisions in terms of risk-return trade-offs or the ability to respond to the awareness of uncertainty.

Our focus in this paper is to study an effective approach to utilize individualized data for the distributional modeling of surgery duration, and to then use this distributional information in the optimization of surgery scheduling. This allows us to account for the uncertainty or variability at a more granular level than before, which turns into subsequent better planning decisions. In particular, we use the quantile regression forest (QRF) pioneered by [96]. The QRF integrates the quantile distribution perspective with the random forest (RF) [26], an ensemble technique that pools together individual decision trees as the base learners. RF and ensemble techniques have shown to be effective in making stable predictions [55], possesses variance-reducing properties [27], and is one of the most popular off-the-shelf machine learning tools. Whereas RF aims to output a mean prediction, QRF, with a very similar procedural design, can make prediction of the whole response distribution (equivalently all the quantiles, hence its name). Compared to traditional linear quantile regression [80], QRF can better capture nonlinear relations among the features and result in less bias, similar to other tree-based bagged learners.

We leverage this powerful tool in surgery duration modeling, with the goal to demonstrate its effectiveness in terms of basic theory and practice. Specifically, we focus on the single-day/block single-OR appointment scheduling. Aligning with the common practice in some surgery departments, e.g. cancer centers, we only consider elective surgeries where the appointments are made a few days prior to the surgery, and hence assume knowledge of the characteristics of the incoming patients (namely the "features" in the predictive model, which can be considered biometric information). Given the predicted distributions of the

surgery durations, the planning problem we consider is to determine the best starting time of each surgery, where the first job always starts at time zero. The trade-off can be described as follows: If the starting time of a subsequent job is too late, then it is more likely that the prior job will finish before the subsequent job arrives, which wastes working time and on a sample path basis and causes the final surgical case to complete at a later time. This is important because jobs at the end of the day incur overtime cost if they complete after the scheduled block time. On the other hand, if the job's start time is too early, then this job is more likely to need to wait, which has undesirable effects that can include increased system congestion, patient dissatisfaction, and negative health outcomes.

The tradeoff above, which arises from the stochasticity of the patient's case durations, can be addressed by three common approaches in optimization. First is the stochastic optimization through Sample Average Approximation (SAA) that approximates an expectation objective, i.e., focuses on the average performance; e.g., [118]. Second is the robust optimization (RO) that focuses on the worst-case performance given support information on the case durations; e.g., [16, 18]. Third is the more recent distributionally robust optimization (DRO) approach based on partially informed probability distributions; e.g., [15, 142]. We show how to incorporate the individualized predictive distribution obtained from QRF into each of these formulations. In the RO and DRO settings in particular, we study uncertainty sets most naturally deduced from QRF, and their tractabilities for the surgery problem by leveraging results in [92]. We also demonstrate how the statistical consistencies of our formulations (with respect to their own target guarantees) follow from the known properties of QRF. We then test these formulations in the context of Memorial Sloan Kettering Cancer Center (MSKCC), using its data. These data contain patient features such as co-morbidity measures, Body Mass Index, demographic information and medical history.

It is worth mentioning that, due to the availability of Electronic Health Records [4, 101], the notion of individualized optimization under uncertainty considered in this paper can be incorporated in various areas of healthcare operational decision making. For instance, our individualized stochastic optimization framework can incorporate operating capacity constraints of hospital resources [73] and chemotherapy infusion treatment sessions scheduling [35] with appropriate adjustments.

We close this introduction with a brief review of related work that brings individualized information into decision-making. The main body of work studies conventional risk minimization or regularization in finding good decisions in areas such as portfolio optimization [10], inventory management [11], online assortment problems [75], medical dosing [14] and the like. In the RO paradigm, [61] used regression models to calibrate uncertainty sets for asset returns in financial contexts. [134] suggested several machine learning approaches to create uncertainty sets. In contrast to the latter, we focus on learning distributions, instead of uncertainty set geometry, which can be incorporated into SAA and DRO and is also more flexible in guaranteeing tractability of the uncertainty sets when using RO.

The remainder of this paper is organized as follows: section 3.2.2 introduces our setting and notations in surgery scheduling. section 3.2.3 reviews QRF, the tool we will use for individualization. section 3.2.4 then demonstrates how we integrate QRF into optimization under the distributional uncertainty of the surgery duration, and presents the corresponding consistency results, in using the three common optimization approaches in the literature. Finally, section 3.2.5 provides an numerical example using MSKCC data to compare our proposed methodologies to the traditional methods.

3.2.2. The Surgery Scheduling Model

This work focuses on the single operating room scheduling problem when the sequence of operations is known. However, the model setting is identical for problems in other areas where the adjustment of time allowances for a sequence of jobs is used for improved operational behavior or cost containment [115, 140]. The model captures the optimal allocation of start times when the surgery duration is uncertain and the number of surgeries is given. The objective is to minimize the the waiting time for the patient and surgery team as well as the overtime for going beyond the scheduled closing time. To simplify the modeling framework, we do not include the idling time in our problem statement. However, it can be easily introduced to the model with a change of notation.

The model focuses on setting the start times of each surgery for a set of n elective surgeries in a given sequence on a given day for a given operating room. We follow similar notations as in [49]. Let T be the scheduled closing time of the day and z_i define the random duration of the ith surgery. The decision maker has to set the surgery time allowance of the ith case, x_i, such that the first starts at time zero, the second at x_1, the third at x_1+x_2, and so on. Assuming that the start time for the first surgery is zero, the start time of any subsequent surgery is scheduled at the sum of the time allowances of all previous surgeries. We also denote by $Z = (z_i)_{i=1,...,n}$ and $X = (x_i)_{i=1,...,n}$ the vectors of surgery durations and surgery time allowances. Based on these assumptions, the waiting time, w_i, and the overtime, l, can be represented as:

$$w_i = \max\{0, w_{i-1} + z_{i-1} - x_{i-1}\}, \qquad i = 2, ..., n \tag{3.4}$$

$$l = \max\{0, w_n + z_n - x_n\} \tag{3.5}$$

We assume that the first surgery always starts on-time and, hence, $w_1 = 0$. For delay in

start of the following surgeries, we incur a unit holding cost for each unit of waiting time (without loss of generality through a rescaling). Additionally, an overtime cost of ϕ per unit time is incurred if we run later than time T. Given the cost function $f(X, Z) = \sum_{i=2}^{n} w_i + \phi l$ and the definitions above, a surgery scheduling model, assuming hypothetically that Z are known, is constructed as:

$$\begin{aligned} \min_{x \in \mathcal{X}} \quad & f(X, Z) \\ s.t. \quad & w_{i+1} = \max\{0, w_i + z_i - x_i\} \qquad i = 1, ..., n-1 \\ & l = \max\{0, w_n + z_n - x_n\} \end{aligned} \tag{3.6}$$

where $\mathcal{X} = \{x_i \geq 0 \ \forall i, \ \sum_{i=1}^{n} x_i \leq T\}$ and $(x_1, ..., x_n)$ are the decision variables. By a monotonicity argument on the objective function in terms of w_i's and l, it is standard to see that (3.6) can be reformulated as

$$\begin{aligned} \min_{\substack{x \in \mathcal{X} \\ w, l}} \quad & f(X, Z) \\ s.t. \quad & w_2 \geq z_1 - x_1 \\ & w_{i+1} \geq w_i + z_i - x_i \qquad i = 2, ..., n-1 \\ & l \geq w_n + z_n - x_n \\ & w_i, l \geq 0 \qquad\qquad\qquad i = 2, ..., n \end{aligned} \tag{3.7}$$

where $(w_2, ..., w_n, l)$ are introduced as auxiliary decision variables. When the Z are stochastic, we can replace the objective of (3.7) as either $E[f(X, Z)]$ or $\{q : P(f(X, Z) \leq q)\}$, where $E[\cdot]$ and $P(\cdot)$ are the expectation and probability taken with respect to Z. The former is an expected value formulation, and the latter is a percentile optimization formulation. Common approaches like SAA and RO provide approximate solutions to these

formulations, as we shall describe in the sequel.

3.2.3. Conditional Distributions and Quantile Regression Forests

To individualize problem (3.7) under stochasticity, we need the distribution of each z_i that is tailored to the characteristic of the patient. Denoting $\xi \in \Xi$ as the patient feature (e.g., gender, age etc.), this means we want to approximate $F(z|\xi)$, the distribution function of the surgery duration given ξ.

One common approach to obtain $F(z|\xi)$ is to use linear quantile regression (LQR) in classical statistics. This estimates the quantile of the distribution conditional on the feature, at a chosen probability level, based on a linear relation between the quantile and the feature [80]. This can be done for different probability levels, leading to estimates of different points of the conditional distribution function. However, for complex and high-dimensional data, the linear assumption in LQR may not hold, and its predictive power is strained.

We therefore consider predictive models based on decision trees, which are flexible to handle nonlinearity. Decision trees are built by partitioning the data into subsets of leaf nodes on the input feature space. Depending on the type of response variables, the partitioning is conducted using criteria of node impurity such as Gini impurity, entropy, variance, etc. [25, 106, 107], which quantify the value or information gain of branching at each node and guide the subsequent steps of the algorithm, until a predefined stopping criterion is met (e.g., maximum tree depth, minimum number of cases in a node, maximum number of nodes in the tree, etc. [25]). Finally, the prediction at a feature value is made by using the mean of the observations within the same leaf node.

Decision trees allow for nonlinearity and thus uncover more complex feature interdependencies than LQR. Moreover, it has been shown to be robust against outliers [25].

The drawback, however, is that they suffer from high variability (over-fitting). Hence, an ensemble of decision trees is normally a preferred choice. In particular, Random Forests (RF) [26] uses multiple trees on different bootstrapped samples of the data, where each of these trees have splitting rules based only on a random subset of the features. The final prediction of RF is made by taking the average of all the decision trees. This averaging, named bootstrap aggregating (or bagging), is known to improve the stability and accuracy of RF compared to single decision trees (e.g., [27, 55]).

Nevertheless, the RF is designed to predict the conditional mean response. The Quantile Regression Forest (QRF) method expands the scope of RF to predict the whole conditional distribution function of the response. The procedure of QRF is similar to RF, but instead of taking the mean of responses within the same leaf as the output of a tree, it takes the empirical distribution of the responses [86]. QRF is shown to perform well in challenging, non-Gaussian situations [96, 135]. More detailed descriptions of RF and QRF are provided in Appendix B.7.

We close this section by stating a consistency result of [96], which shows that QRF recovers the true conditional distribution as the number of observations in the training data increases under appropriate regularity conditions:

Theorem 3.1 (Consistency of QRF; [96])**.** *Under the conditions listed in Appendix B.6, it holds that, for every feature value* $\xi \in \Xi$*,*

$$\sup_{z \in \mathcal{R}} |\hat{F}(z|\xi) - F(z|\xi)| \xrightarrow{p} 0 \quad N \to \infty.$$

where N *denotes the number of i.i.d. observations,* $F(z|\xi)$ *is the conditional cumulative distribution function of the response variable and* $\xrightarrow{p}$ *denotes convergence in probability.*

3.2.4. Individualized Optimization under Uncertainty

This section presents our integration of QRF into different optimization formulations that capture the stochasticity of the surgery duration in an individualized fashion. We first consider an expected value formulation and SAA in section 3.2.4. Then we move to percentile formulation and RO in section 3.2.4, and finally DRO in section 3.2.4.

Expected Value Optimization and Sample Average Approximation

Our first considered formulation is (3.7) but with an expected value objective function

$$E[f(X,Z)|\xi_1,\ldots,\xi_n] = \int f(X,Z) \prod_{i=1}^{n} F(dz_i|\xi_i), \tag{3.8}$$

which minimizes the average overall scheduling cost conditional on the feature ξ_i of each patient, who uses a surgery duration z_i. In this objective, we assume the distributions of all patients are conditionally independent, and we approximate the distribution $F(z|\xi_i)$ by QRF, namely $\hat{F}(z|\xi_i)$. For convenience, we denote $\hat{E}[\cdot|\xi_1,\ldots,\xi_n]$ as the conditional expectation under independent $\hat{F}(z|\xi_i)$'s. Exact computation of the expectation $\hat{E}[f(X,Z)|\xi_1,\ldots,\xi_n]$ in this setting is challenging, not to mention optimization. Thus, we use sample average approximation (SAA) to approximate the problem. In particular, we generate a set of scenarios under $\prod_{i=1}^{n} \hat{F}(z_i|\xi_i)$, denoted $\mathcal{S} = \{s_1, ..., s_k\}$. Every scenario $s \in \mathcal{S}$ is associated with a realization of $Z(s) = (z_i(s))$ and with $w_1(s) = 0$

$$\begin{aligned} w_i(s) &= \max\{0, w_{i-1}(s) + z_{i-1}(s) - x_{i-1}\}, && i = 1, ..., n \\ l(s) &= \max\{0, w_n(s) + z_n(s) - x_n\}. \end{aligned}$$

From (3.7) we can approximate the expected value problem via SAA as

$$\begin{aligned}
\min_{\substack{x\in\mathcal{X}\\ w,l}} \quad & \sum_{k=1}^{K}\frac{1}{K}\Big(\sum_{i=2}^{n} w_i(s_k)+\phi l(s_k)\Big) && && (3.9)\\
s.t. \quad & w_{i+1}(s_k)\geq w_i(s_k)+z_i(s_k)-x_i && i=1,...,n-1, k=1,...,K\\
& l(s_k)\geq w_n(s_k)+z_n(s_k)-x_n && k=1,...,K\\
& w_i(s_k), l(s_k)\geq 0 && i=2,...,n, k=1,...,K.
\end{aligned}$$

Formulation (3.9) naturally combines the distributional prediction of QRF into the expected value minimization. This is in contrast to previous approaches that typically treat all patients identically and all the z_i are generated in an i.i.d. fashion.

Under standard conditions, the SAA problem (3.9) has a solution and optimal value that converge to those of (3.8), as we formally state next.

Theorem 3.2. *Let*

$$X^* \in \operatorname*{argmin}_{X\in\mathcal{X}} E[f(X,Z)|\xi_1,\ldots,\xi_n] \tag{3.10}$$

and $\tilde{X}^*$ be an optimal solution to (3.9)*. Assume that z and ξ satisfies the conditions in Appendix B.6, we have $H(\tilde{X}^*)\xrightarrow{p} H(X^*)$ as $K,N\to\infty$, where N is the observation size in building the QRF and K is the scenario size.*

Robust Optimization

Robust optimization (RO) provides an alternate paradigm to account for the uncertainty in the surgery duration. In the RO paradigm, we replace the stochasticity with a so-called uncertainty set or ambiguity set, which is a deterministic set that, intuitively, captures the likely realization of the surgery duration (other interpretations are possible, e.g., [16], P.33 discussion point B).

Formulation (3.7) has resemblance with single-server queues, for which [12] has introduced an RO approach to estimate relevant quantities. Following their framework, we introduce $\underline{\Gamma_k}$ and $\overline{\Gamma_k}$ and assert that the surgery times belong to the uncertainty set

$$\mathcal{U} = \left\{ (z_1, z_2, ..., z_n) \;\middle|\; \underline{\Gamma_{k,i}} \leq \sum_{j=k}^{i-1} z_j \leq \overline{\Gamma_{k,i}} \;, \quad \forall k < i \leq n \right\}, \tag{3.11}$$

The reason why we consider a set on $\sum_{j=k}^{i-1} z_j$, instead of other possible candidates (e.g., merely z_k themselves), is that the objective function explicitly depends on the former, which facilitates solution to the RO.

RO considers

$$\min_{x \in \mathcal{X}} \max_{z \in \mathcal{U}} \left(\sum_{i=2}^{n} w_i + \phi l \right) \tag{3.12}$$

where w_i and l satisfy (3.4) and (3.5). Note that the waiting time of the ith patient can be expressed as

$$w_i = \max\{w_{i-1} + z_{i-1} - x_{i-1}, 0\} = \max_{1 \leq k < i} \left(\sum_{j=k}^{i-1} z_j - x_j, 0 \right). \tag{3.13}$$

and it is also trivial to see that $l = w_{n+1}$. Putting these into (3.12) gives

$$\min_{x \in \mathcal{X}} \max_{z \in \mathcal{U}} \left(\sum_{i=2}^{n} \max_{1 \leq k < i} \left(\sum_{j=k}^{i-1} z_j - x_j, 0 \right) + \phi \max_{1 \leq k < i} \left(\sum_{j=k}^{n} z_j - x_j, 0 \right) \right). \tag{3.14}$$

We then switch the order of maximizations to achieve

$$\min_{x \in \mathcal{X}} \left(\sum_{i=2}^{n} \max_{1 \leq k < i} \max_{z \in \mathcal{U}} \left(\sum_{j=k}^{i-1} z_j - x_j, 0 \right) + \phi \max_{1 \leq k < i} \max_{z \in \mathcal{U}} \left(\sum_{j=k}^{n} z_j - x_j, 0 \right) \right). \tag{3.15}$$

The innermost maximization can be easily seen to be attained at the upper bounds imposed

in $\mathcal{U}$, so that (3.15) is equivalent to

$$\min_{x\in\mathcal{X}} \left(\sum_{i=2}^{n} \max_{1\le k<i} \left(\overline{\Gamma}_{k,i-1} - \sum_{j=k}^{i-1} x_j, 0 \right) + \phi \max_{1\le k<i} \left(\overline{\Gamma}_{k,n} - \sum_{j=k}^{n} x_j, 0 \right) \right). \tag{3.16}$$

Now, denoting $Q_i = \max_{1\le k<i} \left(\overline{\Gamma}_{k,i-1} - \sum_{j=k}^{i-1} x_j, 0 \right)$, then (3.16) can be reformulated into the following linear program

$$\begin{aligned} \min_{x\in\mathcal{X}} \quad & \sum_{i=2}^{n} Q_i + \phi Q_{n+1} \\ s.t. \quad & Q_i \ge \overline{\Gamma}_{k,i-1} - \sum_{j=k}^{i-1} x_j \qquad i = 1, ..., n+1 \text{ and } 1 \le k < i \\ & Q_i \ge 0 \qquad\qquad\qquad\qquad i = 2, ..., n+1. \end{aligned} \tag{3.17}$$

The question remains how to calibrate $\overline{\Gamma}_{k,i}$'s. Using the idea of data-driven RO (e.g., [19]), one can set $\overline{\Gamma}_{k,i}$'s so that

$$\hat{P}\left(\underline{\Gamma_{k,i}} \le \sum_{j=k}^{i-1} z_j \le \overline{\Gamma_{k,i}} \ , \quad \forall k < i \le n \middle| \xi_1, \ldots, \xi_n \right) \ge 1 - \delta,$$

where $\hat{P}(\cdot|\xi_1, \ldots, \xi_n)$ denotes the probability under $\prod_{i=1}^{n} \hat{F}(z_i|\xi_i)$. In particular, we can find Γ such that

$$\hat{P}\left(\left| \frac{\sum_{j=k}^{i-1} z_j - \sum_{j=k}^{i-1} \mu_j}{\sqrt{\sum_{j=k}^{i-1} \sigma_j^2}} \right| \le \Gamma \ , \quad \forall k < i \le n \middle| \xi_1, \ldots, \xi_n \right) \ge 1 - \delta \tag{3.18}$$

where μ_j and σ_j^2 are the mean and variance from $\hat{F}(\cdot|\xi_j)$. The left hand side of the expression inside the probability in (3.18) is a centered and normalized version of $\sum_{j=1}^{i-1} z_j$ that appears often in the central limit theorem. The choice of Γ in (3.18) then implies that

one can choose

$$\underline{\Gamma}_{k,i} = \sum_{j=k}^{i-1} \mu_j - \Gamma \sqrt{\sum_{j=k}^{i-1} \sigma_j^2}$$

and

$$\overline{\Gamma}_{k,i} = \sum_{j=k}^{i-1} \mu_j + \Gamma \sqrt{\sum_{j=k}^{i-1} \sigma_j^2} \tag{3.19}$$

Now, to find Γ that satisfies (3.18), we can use the quantile estimation approach (suggested by, e.g., [69]). Simulate, say K, i.i.d. copies of Z, and for each Z one can calculate

$$\max_{k<i\leq n} \left| \frac{\sum_{j=k}^{i-1} z_j - \sum_{j=k}^{i-1} \mu_j}{\sqrt{\sum_{j=k}^{i-1} \sigma_j^2}} \right| \tag{3.20}$$

Find the $\lfloor K(1-\delta) \rfloor$-th order statistic of the K copies of (3.20), call it $\tilde{\Gamma}$. This will be a consistent estimate of Γ that satisfies (3.18).

Using the above formulations and procedure gives the following:

Theorem 3.3. *Assume that z and ξ satisfies the conditions in Appendix B.6. Let $\tilde{H}$ be the optimal value of formulation* (3.17)*, with $\overline{\Gamma}_{k,i}$ calibrated from* (3.19) *with Γ approximated by $\tilde{\Gamma}$, the $\lfloor K(1-\delta) \rfloor$-th order statistic of* (3.20) *derived from the K samples of Z. Let H^* be the optimal value for $\min_{x\in\mathcal{X}} \min\{q : P\left(f(X,Z) \leq q | \xi_1, \ldots, \xi_n\right) \geq 1-\delta\}$. Then*

$$\liminf_{N\to\infty, K\to\infty} \tilde{H} \geq H^*$$

Note that the optimization $\min_{x\in\mathcal{X}} \min\{q : P\left(f(X,Z) \leq q | \xi_1, \ldots, \xi_n\right) \geq 1-\delta\}$ means minimizing the true $(1-\delta)$-quantile of $f(X,Z)$. Theorem 3.3 stipulates that the RO formulation can be viewed as a conservative approximation of this optimization.

Distributionally Robust Optimization

We next consider distributionally robust optimization (DRO). This approach targets expected value objective function, like in the case of SAA, but under only partial information of the distributions. More specifically, it optimizes the worst-case expected value among all distributions that are in an uncertainty set or an ambiguity set that contains the partial information. In our scheme, we assume the quantiles of each patient's surgery duration distribution at a list of given probability levels are known. These information pieces can be drawn from the QRF, achieving individualization.

The motivation for using DRO is that its solution tends to be more robust. In our circumstance, imposing enough quantile information means we know the distribution of each patient's surgery duration distribution, but not assuming, e.g., any dependency structure among the patients. The solution from DRO is thus robust against such type of hidden stochasticity that is not revealed by the individualization, which focuses only on the prediction for each patient.

More concretely, for each patient i, we consider a set of probability $r_{ij}, j = 1, \ldots, m$, and we denote by q_{ij} the r_{ij}-th quantile of the duration distribution given ξ_i, which can be inspected from the QRF. For convenience, we set $r_{im} = 1$. We consider the uncertainty set

$$\mathcal{U} = \{P \in \mathcal{P} : P(z_i \leq q_{ij}) = r_{ij},\ i = 1, \ldots, n, j = 1, \ldots, m\} \tag{3.21}$$

where $\mathcal{P}$ is the set of all probability distributions supported on $\mathbb{R}^n_+$. Here, r_{ij} and q_{ij} are calibrated so that $\hat{F}(q_{ij}|\xi_i) = r_{ij}$, for all i, j, which indicates the information that the marginal r_{ij}-th quantiles for patient i are known to be q_{ij} under the QRF. This, in particular, includes the information that the largest possible value of $\hat{F}(z|\xi_i)$ is estimated to be q_{im}.

We seek to solve

$$\min_{x\in\mathcal{X}} \max_{P\in\mathcal{U}} E_P\big[f(X,Z)\big]. \tag{3.22}$$

We approach (3.22) using the technique in [92], which first replace $f(X,Z)$ as the optimal value of a linear program (LP) given by

$$\begin{aligned} \max_{y} \quad & \sum_{i=1}^{n}(z_i - x_i)y_i \\ s.t. \quad & y_i - y_{i-1} \geq -1 \qquad 2 \leq i \leq n \\ & y_n \leq \phi \\ & y_i \geq 0 \qquad \forall i = 1, ..., n. \end{aligned} \tag{3.23}$$

This can be derived from the definition of w_i and l in (3.4) and (3.5), and considering the dual of the resulting LP, where $y_1, \ldots, y_n$ are the dual variables. Then problem (3.22) can be rewritten as

$$\min_{x\in\mathcal{X}} \max_{P\in\mathcal{U}} E_P\big[\max_{y\in\Omega} \sum_{i=1}^{n}(z_i - x_i)y_i\big], \tag{3.24}$$

where Ω refers to the constraint set in (3.23). We now focus on the outer maximization problem, $\max_{P\in\mathcal{U}} E_P\big[\max_{y\in\Omega}\sum_{i=1}^{n}(z_i - x_i)y_i\big]$. The following lemma transforms it into a more manageable form:

Lemma 3.1. *The dual representation of* $\max_{P\in\mathcal{U}} E_P\big[\max_{y\in\Omega}\sum_{i=1}^{n}(z_i - x_i)y_i\big]$ *is equivalent to*

$$\min_{\rho} \max_{y\in\Omega} \sum_{i=1}^{n} \max_{j=1,\ldots,m} \left\{(q_{ij} - x_i)y_i - \sum_{k\geq j}^{m} \rho_{ik}\right\} + \sum_{i=1}^{n}\sum_{j=1}^{m-1} r_{ij}\rho_{ij} + \sum_{i=1}^{n} \rho_{im}, \tag{3.25}$$

where $\rho_{11}, \ldots, \rho_{nm}$ *are the dual variables corresponding to the quantile constraints.*

By Lemma 3.1, problem (3.22) can be represented by

$$\min_{x\in\mathcal{X},\rho}\max_{y\in\Omega}\sum_{i=1}^{n}\max_{j=1,...,m}\left\{(q_{ij}-x_i)y_i-\sum_{k\geq j}^{m}\rho_{ik}\right\}+\sum_{i=1}^{n}\sum_{j=1}^{m-1}r_{ij}\rho_{ij}+\sum_{i=1}^{n}\rho_{im}. \tag{3.26}$$

We want to transform (3.26) to a linear program. Note that $\max_{y\in\Omega}\sum_{i=1}^{n}\max_{j=1,...,m}\left\{(q_{ij}-x_i)y_i-\sum_{k\geq j}^{m}\rho_{ik}\right\}$ is a convex maximization problem over the polyhedron Ω, and it suffices to consider its extreme points for an optimal solution. This allows us to follow the method discussed in Proposition 2 of [92] to reduce to solving an LP relaxation of an integer program using the following proposition.

Proposition 3.1. *Problem* (3.22) *can be reformulated as the following linear program*

$$\begin{aligned}
\min_{x,\rho,\lambda}\quad & \sum_{i=1}^{n}\lambda_i+\sum_{i=1}^{n}\sum_{j=1}^{m-1}r_{ij}\rho_{ij}+\sum_{i=1}^{n}\rho_{im} & \\
s.t.\quad & \mu_{io}\geq\left\{q_{ij}\pi_{io}-\sum_{k\geq j}^{m}\rho_{ik}\right\} & 1\leq i\leq n, 1\leq i\leq o\leq n+1, 1\leq j\leq m\\
& \sum_{i=g}^{min\{o,n\}}\mu_{io}\leq\sum_{i=g}^{min\{o,n\}}x_i\pi_{io}+\lambda_i & 1\leq g\leq n, g\leq o\leq n+1\\
& \sum_{i=1}^{n}x_i\leq T & \\
& x_i\geq 0 & \forall i\in\{1,...,n\},
\end{aligned} \tag{3.27}$$

where $(\lambda_1,...,\lambda_n)$ *are the dual variables,* $\pi_{io}=y_i$ *for all* $i\in[g,o]$ *and any* $g\leq o$. *Also,* $\mu_{io}=\max_{j=1,...,m}\{q_{ij}\pi_{io}-\sum_{k\geq j}^{m}\rho_{ik}\}$ *for* $o=1,...,n+1$.

The following shows that the DRO under the quantile information drawn from the QRF gives an asymptotic bound for the true expected value:

Theorem 3.4. *Assume that* z *and* ξ *satisfies the conditions in Appendix B.6. Let* H^* *be*

the optimal value of (3.22)*, and* $\hat{H}$ *be the optimal value of* (3.27)*. Then for any* $\epsilon > 0$*, we have*

$$P\left(H^* \leq \hat{H} + \epsilon\right) \to 1 \tag{3.28}$$

as $N \to \infty$*, where* $\mathcal{U}$ *is the uncertainty set defined in* (3.21) *and* N *is the data size in constructing the QRF.*

Note that Theorem 3.4 does not require strong duality of the dual formulation in (3.25), because it involves only an upper bound guarantee in (3.28). Moreover, unlike SAA and RO discussed before, our DRO reformulation does not require running simulation from the QRF.

3.2.5. Numerical Results and Discussion

We test our approach against some benchmark procedures. Our goal is to give insight into how well our three optimization approaches can harness the individualized case duration distributions to improve the decision quality for the surgery scheduling problem.

To implement our approach, we construct a QRF over the case duration data. We use 80% of the available data for training and the remaining 20% for testing the QRF when integrated with SAA, DO, and DRO. A detailed description of the data is given in §3.2.5. This QRF will provide the individualized distributions, i.e., conditional distributions given the patients' features, to feed into our optimization formulations as presented in §3.2.4.

To evaluate our performance, we compare with several commonly used approaches. Firstly, we fit a log-normal distribution on all case durations in the training set (L-Norm). It is well known that surgeons can vary significantly in surgical distribution for the same type of surgery. For this reason, we Secondly stratified the data by the primary surgeons and fit a log-normal distribution on all case durations in the training set (L-Norm-PS). The

log-normal distribution was selected because it delivered the best fit compared to other well-known distributions. For this we used the Kolmogorov-Smirnov test in R to identify the quality of fit for several well-known distributions including Normal, Weibull, Gamma, etc. Third, we use a linear quantile regression constructed from the training set (LQR). For each of the optimization approaches (SAA, RO and DRO), we compare against the performances of these distribution modeling methods as well as QRF using the patient data in the test set.

Data

We consider the surgeries performed at Memorial Sloan Kettering Cancer Center (MSKCC) within the last six years (2010 - 2016). MSKCC is one of the leading cancer centers in the world which is constantly improving cancer treatment process by applying state-of-the-art scientific methodologies in both medical and operational aspects of its care delivery system. Its surgical operations take place across thirteen different services and in a total of forty different operating rooms. The overall dataset contains a total of $129,742$ distinct surgeries in all services throughout the study horizon, and it included all inpatient and outpatient surgeries. For illustration, we present only the results for Urology (URO) service for which we have $23,176$ observations. Table 3.9 provides summary information regarding the number of URO surgeries, the number of unique primary surgeons, the number and proportion of inpatient surgeries and the historical durations of surgeries in URO service.

Table 3.9: Characterization of operational data for URO service in the hospital

Service	Surgeries	Surgeons	Inpatient (%)	Case Duration (SD)
URO	23,176	49	9,428 (40.7)	153.5 (129.2)

The features collected in the historical data can be classified into the main groups of

Table 3.10: Demographics of the patient data in URO service

			Race, Number of Observations (%)			
Service	Age (SD)	Female (%)	White	Black	Asian	Other
URO	63.1 (13.6)	5196 (22.4)	19,544 (84.3)	1,466 (6.3)	799 (3.4)	1,367 (5.9)

patient, surgery team, equipment, operational and temporal data. On the patient side, we include age, gender, race, BMI, treatment history and comorbidity measures. For example, Table 3.10 summarizes key demographics of patient data we use as features. Moreover, we incorporate data on Relative Value Units (RVU), equipment, surgeons, operational specifications of surgeries, etc. A detailed list of features used in the final predictive model is given in table 3.11, which is also used in [99] but for a different purpose.

Test of Performance

We compare the results for the URO surgery schedules with number of surgeries varying from three to six. Specifically, in the test set, the number of cases performed in a day varies as follows: 754 schedules with three surgeries, 281 with four surgeries, 303 with five surgeries and 287 with six surgeries. We assume equal penalties for waiting and overtime. Also, we consider T to be equal to the sum of the actual surgery durations in the schedule.

We start with the SAA framework discussed in section 3.2.4. For schedule associated with each surgery numbers, we implement SAA by drawing 5000 samples from the surgery duration distributions that we derive from the L-Norm, L-Norm-PS, LQR and QRF for all the surgeries in that schedule. Figure 3.5 shows the distribution of waiting , overtime and overall costs for each of the four models, broken down by the number of surgeries each schedule has in the test set. Our individualized approach using QRF appears to outperform the other three models for all four types of schedules, with the outperformance more significant for the schedules with bigger numbers of surgeries. More concretely, Table

Table 3.11: Features used for training predictive models

Category	Features
Hospital	The primary surgeon, experience of the primary surgeon in years, number of publications of the primary surgeon, surgery room, number of panels in the surgical case, the sum of all RVU measures for the case, the mean of all RVU measures for the case, the maximum of all RVU measures for the case, the minimum of all RVU measures for the case, number of procedures for the surgical case, number of robotic procedures, the mean duration of the last 5 similar cases based on the primary procedure, the number of times the surgeon had a similar surgery in the 30 days, the number of times the surgeon had a surgery in the 30 days, the number of times the surgeon had a similar surgery in the past, the number of times the surgeon a surgery in the past, whether or not a particular equipment was needed for the surgery, the number of required equipment
Patient	Comorbidity measures such as obesity, depression, etc., Body Mass Index (BMI), weight, race, gender, age, whether the patient is inpatient or outpatient, the number of days spent in the hospital prior to the surgery, the number of times patient underwent chemotherapy, the number of times patient underwent radiation therapy, the number of times patient underwent surgeries
Operational	The number of days between the day of surgery and the day the surgery was scheduled, the number of cases assigned to the surgeon on the day of surgery, the number of cases assigned to the surgery room on the day of surgery, the number of cases scheduled on the day of surgery, the sequence number of the surgery on the surgeon's schedule, the sequence number of the surgery in the surgery room
Temporal	The weekday of the surgery, the month of the surgery, the year of the surgery

3.12 shows our percentage improvements in three different performance metrics (mean, 75th and 95th percentiles) of the overall cost. Across all settings, QRF outperforms other

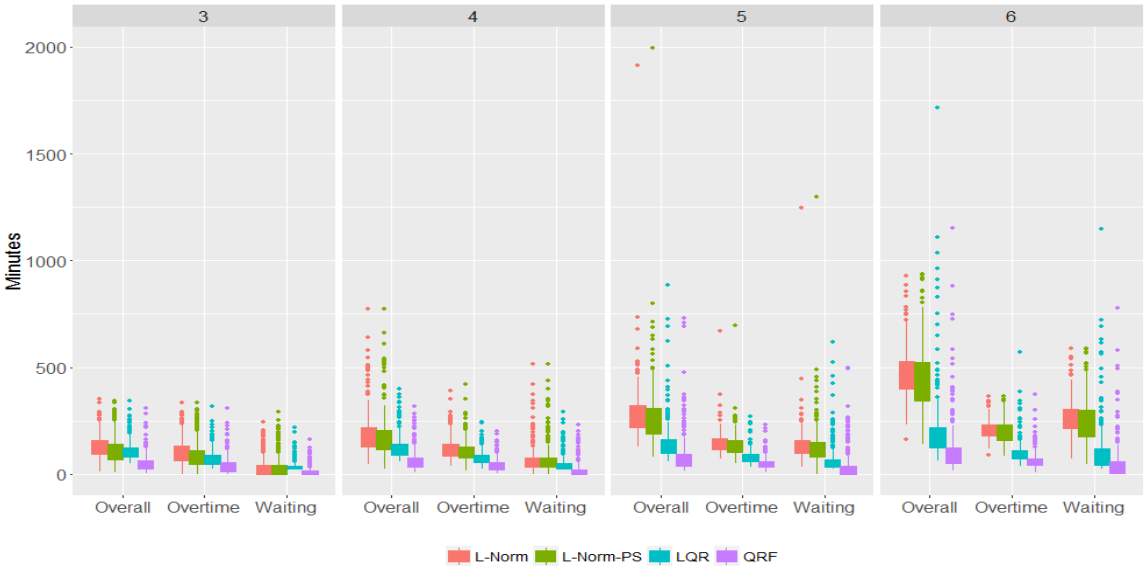

Figure 3.5: The Box-plot for waiting, overtime and overall cost for the schedules with 3, 4, 5 and 6 surgeries in the test set based on the SAA method using the L-Norm, L-Norm-PS, LQR and QRF models

methods in all surgery groups by an overall average of 60.2%, 59.6% and 41.2% in mean, 75th and 95th percentiles respectively.

Table 3.12: Percentage improvement of the overall cost under different performance metrics between QRF model and the L-Norm, L-Norm-PS, LQR models based on the Individualized Stochastic Optimization framework

		Number of Surgeries			
Metric	Model	3	4	5	6
Mean	L-Norm	60.8	66.8	69.1	75.6
	L-Norm-PS	55.4	64.5	67.2	74.0
	LQR	53.2	49.6	41.8	44.6
75th percentile	L-Norm	58.5	65.1	70.9	76.7
	L-Norm-PS	53.6	63.6	70.0	76.6
	LQR	47.7	47.2	42.0	43.4
95th percentile	L-Norm	51.6	56.7	36.0	49.3
	L-Norm-PS	53.3	56.5	39.2	51.7
	LQR	35.2	30.2	6.6	28.8

Table 3.13: Percentage improvement of the overall cost under different performance metrics between QRF model and the L-Norm, L-Norm-PS, LQR models based on the Individualized Robust Optimization framework

		Number of Surgeries			
Metric	Model	3	4	5	6
Mean	L-Norm	55.7	52.8	59.0	68.9
	L-Norm-PS	44.1	45.0	57.5	68.4
	LQR	27.4	30.0	23.2	18.5
75th percentile	L-Norm	58.4	50.6	58.9	68.0
	L-Norm-PS	43.0	42.2	57.3	67.7
	LQR	20.1	24.8	17.0	17.1
95th percentile	L-Norm	49.6	51.3	46.7	55.3
	L-Norm-PS	40.3	32.4	42.8	55.0
	LQR	9.6	17.1	3.3	5.5

Next we consider the RO framework discussed in section 3.2.4. For each schedule, we implement RO using uncertainty sets based on (3.11) derived from L-Norm, L-Norm-PS,

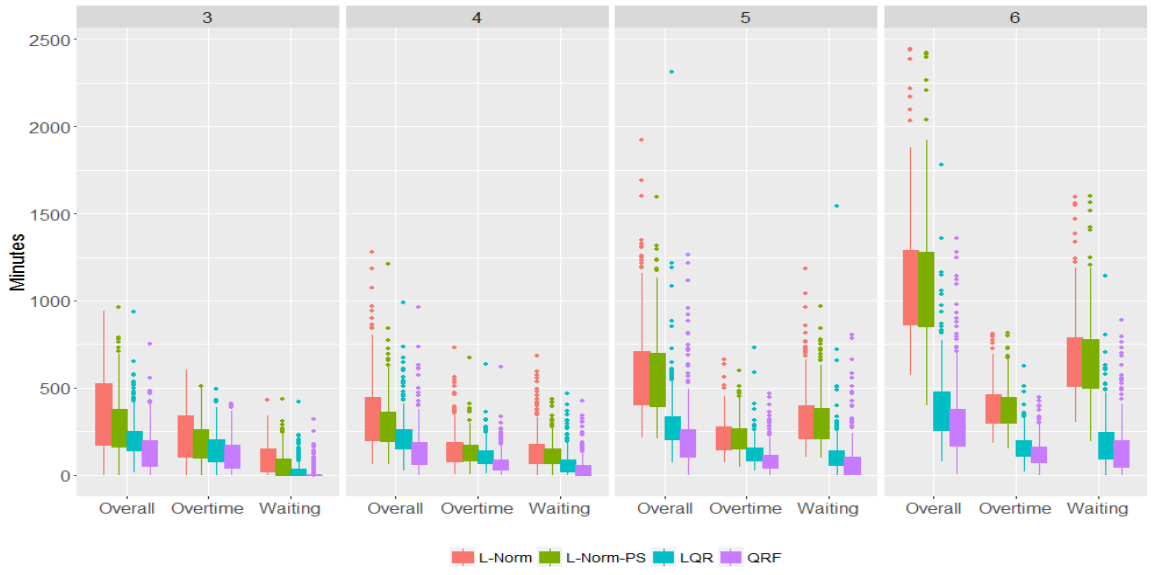

Figure 3.6: The Box-plot for waiting, overtime and overall cost for the schedules with 3, 4, 5 and 6 surgeries in the test set based on the robust optimization framework using the L-Norm, L-Norm-PS, LQR and QRF models

Table 3.14: Percentage improvement of the overall cost under different performance metrics between QRF model and the L-Norm, L-Norm-PS, LQR models based on the Individualized Distributionally Robust Optimization framework

		Number of Surgeries			
Metric	Model	3	4	5	6
Mean	L-Norm	58.8	67.2	67.0	73.6
	L-Norm-PS	52.9	65.1	65.9	71.4
	LQR	50.3	52.3	40.1	40.4
75th percentile	L-Norm	55.4	65.7	67.3	74.0
	L-Norm-PS	49.8	65.7	67.4	73.4
	LQR	44.2	48.7	41.1	41.2
95th percentile	L-Norm	47.1	51.2	51.6	62.7
	L-Norm-PS	49.1	48.8	50.9	64.5
	LQR	33.6	30.2	28.2	31.0

LQR and QRF. Figure 3.6 depicts the distributions of waiting , overtime and overall costs for each of the four models grouped by the number of surgeries each schedule has in the

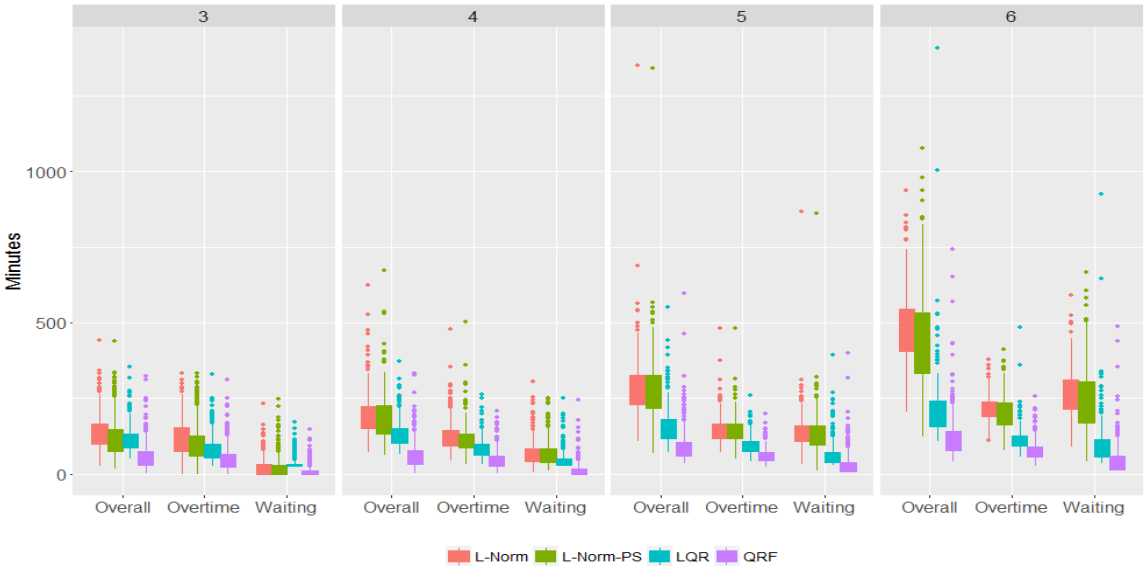

Figure 3.7: The Box-plot for waiting, overtime and overall cost for the schedules with 3, 4, 5 and 6 surgeries in the test set based on the distributionally robust optimization framework using the L-Norm, L-Norm-PS, LQR and QRF models

test set. Our QRF again performs consistently better than the other three models. Table 3.13 indicates the percentage improvement of QRF over the other three methods. Here, however, LQR is a close competitor, with on average 24.7%, 19.7% and only 8.8% less effective than QRF in mean, 75th and 95th percentiles of the overall cost respectively.

Finally, we consider the DRO framework in section 3.2.4. We implement DRO with constraints dictated by percentiles $\{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8,\ 0.9, 0.95\}$ of the surgery duration distributions derived from L-Norm, L-Norm-PS, LQR and QRF. Figure 3.7 once again shows improvement of QRF over the other three models regarding the distributions of waiting , overtime and overall costs. In particular, Table 3.14 shows percentage improvements across all settings of 58.7%, 57.8% and 45.7% in mean, 75th and 95th percentiles of the overall cost respectively.

3.2.6. Conclusion

Surgery scheduling problems involve immense distributional uncertainty. Previous studies consider capturing the distributional information through aggregate information, and individualization is limited to only mean response prediction. In this paper, we studied a framework to individualize distributions of surgery durations by using QRF to generate conditional predictive distributions. We showed how to incorporate predictions from QRF into three common approaches in optimization under uncertainty, namely SAA, RO and DRO. We studied the reformulations and consistent statistical guarantees for each approach. We demonstrated through MSKCC data how our individualization of surgery duration distributions outperformed several established approaches without individualization or with a less powerful prediction tool.

Chapter 4.

Sequential Learning Under Probabilistic Constraints: Feasibility of Digital Marketing Campaigns

4.1. Introduction

Most of the literature in stochastic online learning focuses on performances measured by optimality achievement. Common examples include the minimization of cumulative regret in the multi-arm bandit setting (e.g., [8, 82]), best arm selection (e.g., [7]) and the closely related ranking and selection (e.g., [23]) in the simulation literature. In many situations, however, the uncertainty or the stochasticity appears not only in the objective function, but also the constraints of the problem whose feasibility can be of utmost importance. The focus of this paper is to design sequential methodologies that maintain probabilistic feasibility requirements with rigorous statistical guarantees.

Our study is motivated from a rich set of problems where "budgets" or "resources" are

limited for various operational or commercial reasons, and these constraints are in a sense "soft", i.e., the capacities placed on these constraints, while preferred to be satisfied, are allowed to be violated with a small probability. Such consideration is common among applications. For example, in online advertisement encountered by our co-authors while optimizing ad spend for large advertisers, the task involves sequentially picking items (e.g., keywords, targets) to maximize revenues, while adhering to a specified marketing budget for a duration. The marketer in general expects to meet the budget goals. However, if occasionally the budget is exceeded the campaign is still acceptable as long as the revenue performance is sustained. Other similar settings include clinical trials, where the costs of competing treatments are substantial and noisy, and over-budget is undesired but sometimes allowable. Whereas past work in stochastic sequential learning has focused on rewards (with hard constraints if needed), this paper provides the first study on a class of problems that not just include the rewards but also stochastic constraints that need to be satisfied with high probability.

Our framework can be viewed as a sequential problem under so-called *probabilistic* or *chance constraints* [105], which has been widely used in stochastic programming under limited and uncertain resources (e.g., [84, 119]). A generic representation of a chance constraint is

$$P(f(x,\xi) \in \mathcal{A}) \geq 1-\alpha \tag{4.1}$$

where $x \in \mathbb{R}^d$ is a decision variable, $\xi \in \Xi$ is a random variable (or vector) under P that controls the stochasticity, $f : \mathbb{R}^d \times \Xi \to \mathcal{Y}$ is a known function and $\mathcal{A} \subset \mathcal{Y}$. For instance, $f(x,\xi) \in \mathcal{A}$ can represent a set of inequalities so that $\mathcal{Y} = \mathbb{R}^m$ for some m and $\mathcal{A} = \{y \in \mathbb{R}^m : y \leq 0\}$. The condition $f(x,\xi)$ is sometimes called the safety condition. The given parameter α is the tolerance level that represents the allowable probabilistic violation of the safety condition.

In the sequential setting, x would denote a sequence of decisions. The safety condition could include individual requirements on all or some stages. In many applications of interest, P needs to be learned as complete distributional knowledge on ξ is not available. Along the vein of conventional online problems that focus on optimality, at each stage we may observe some components of ξ so that we can update our belief on P.

As our main methodological contribution, we analyze an online strategy that provides guarantees on (4.1) with a high confidence, under a statistical framework that we shall describe. On a high level, it means we can guarantee, with our proposed policy, that

$$P(P(f(x,\xi) \in \mathcal{A}) \geq 1-\alpha) \geq 1-\beta \tag{4.2}$$

where the outer probability now refers to the randomness of x induced by the sequential observations, and $1-\beta$ is a confidence level (90% for instance). Our methodology is based on a combination of two ideas. First is a Bayesian extension of the so-called scenario generation or constraint sampling [31, 47] approach in approximating chance-constrained optimization problems. This approach replaces the unknown or difficult chance constraint with a collection of sampled constraints that come from data or from numerical simulation. Viewing such an approach in a Bayesian manner allows it to be blended naturally into popularly used online learning algorithms such as Thompson sampling [1, 114] that also operates via Bayesian updating. Second, by capitalizing results on scenario generation in the static setting, we can derive the precise number of samples required at each stage of the sequential process such that (4.2) holds throughout the horizon. As far as we know, our formulation and analysis of chance-constraint guarantees in an online setting is new to the literature.

After presenting our theoretical investigation on feasibility guarantees, we illustrate the integration of our scheme into a variant of Thompson sampling in an online advertisement

setting. We then numerically demonstrate how this chance-constrained Thompson sampling performs competitively, in achieving feasibility but also maintaining good objective values.

4.2. Related Work

The earliest work in chance constraints dated back to [41] and [100]. Exact solution techniques for such problems are notoriously difficult due to non-convexity, and are only available in few instances even when P is known; e.g., [81]. Several lines of approximation methodologies have been proposed. A conventional method is to use so-called safe convex approximation that replaces the chance constraint with more conservative convex constraints [17]. [110, 111] used policy trees and confidence interval construction to obtain the so-called (α, ϑ)-solution. Scenario generation [29, 33], which we leverage on in this work, uses sampled constraints to populate the feasible region. This approach has several extensions, such as sampling-and-discarding [32] and multi-phase schemes [28, 34, 38], and relates to sample average approximation [89]. Other data-driven methods include distributionally robust optimization [30, 150] and data-driven robust optimization [19].

Our work focuses on chance-constrained problem in an online fashion, under the broad umbrella of sequential decision-making. In the later part of this paper, we demonstrate our proposed strategy in a variant of the stochastic multi-arm bandit problem [8] used in addressing the well-known exploration-exploitation tradeoff. In budgeted bandits, [54] consider the presence of random costs and an overall budget, where learning and revenue accumulation stops when the budget runs out. [144] study Thompson sampling for a similar setting; in this work we integrate our strategy into Thompson sampling, especially the one considered in [58] motivated from network revenue management. Other related work include those in the framework of "bandits with knapsacks" [9, 21, 132] that have

been applied in pricing and supply chain management ([141]) and healthcare ([138]). The works closest to our online advertisement example are [133] and [3] that study the problem of item bidding under a budget, but they do not consider probabilistic violation of the constraints that we focus on.

4.3. Chance Constrained Online Learning

We consider a sequential problem with horizon T

$$\begin{aligned} \max_{x_1,\ldots,x_T} \quad & h(x_1,\ldots,x_T) \\ \text{subject to} \quad & P(f_t(x_{1:t},\xi_{1:t}) \in \mathcal{A}_t | \mathcal{F}_{t-1}) \geq 1-\alpha \ \forall t \in \mathcal{S}, \end{aligned} \tag{4.3}$$

where $\xi_t \in \Xi_t$ are random variables assumed independent among the steps t, $x_1,\ldots,x_T$ are the sequential decisions in which x_{t+1} depends on the past observations of $\xi_{1:t} = (\xi_1,\ldots,\xi_t)$ and past decisions $x_{1:t} = (x_1,\ldots,x_t)$, and $\mathcal{S} \subset \{1,\ldots,T\}$ is a given set. The function f_t and the set $\mathcal{A}_t$ can depend on the time step t. For convenience, let $\mathcal{F}_t = \{\xi_{1:t}, x_{1:t}\}$ be the information up to time t. In each probability P in (4.3), the function $f_t(x_{1:t},\xi_{1:t})$ can be expressed as $f_t(x_{1:(t-1)}, x_t, \xi_{1:(t-1)}, \xi_t)$ where $x_{1:(t-1)}$ and $\xi_{1:(t-1)}$ belong to the past information $\mathcal{F}_{t-1}$.

Consistent with the introduction, α is a tolerance parameter on the violation of the safety condition. This parameter is assumed constant across t for convenience, but our analysis can be easily adapted to the case where it varies. Note that $\mathcal{S}$ determines how many chance constraints need to be maintained throughout the horizon. For example, $\mathcal{S} = \{1,\ldots,T\}$ means there is a budget requirement for each step, and $\mathcal{S} = \{T\}$ means there is only one overall budget requirement across the whole horizon.

4.3.1. Scenario Generation For Static Problems

We first discuss a well-studied approach to approximate a static version of (4.3). Suppose $T = 1$. In this setting we can simplify notation and write the formulation as

$$\begin{aligned} \max_{x} \quad & h(x) \\ \text{subject to} \quad & P(f(x,\xi) \in \mathcal{A}) \geq 1 - \alpha \end{aligned} \tag{4.4}$$

Suppose we can simulate or collect data for ξ to obtain, say, i.i.d. $\xi^1, \ldots, \xi^N$. We consider replacing the constraint in (4.4) by sampled constraints, so that the optimization program becomes

$$\begin{aligned} \max_{x} \quad & h(x) \\ \text{subject to} \quad & f(x,\xi^n) \in \mathcal{A}, \ \forall\ n = 1, \ldots, N \end{aligned} \tag{4.5}$$

We call (4.5) a sampled program, which serves as a reasonable approximation to (4.4) when N is large. However, since ξ^n's are randomly generated, the solution obtained from (4.5) is subject to statistical noise and cannot be guaranteed feasible for (4.4). The following celebrated result from [29, 33] gives the sample size needed to guarantee feasibility for (4.4) with a certain confidence by solving (4.5).

Condition 1. *An optimization program is said to be in class $\mathcal{R}$ if: 1) It is feasible and the feasible region has a non-empty interior; 2) Its optimal solution exists and is unique.*

Theorem 4.1. *(Adopted from Theorem 2.4, [33]) Suppose for each ξ, $f(x,\xi) \in \mathcal{A}$ is a convex set in $x \in \mathbb{R}^d$, and h is concave. Suppose also that any instance of* (4.5) *belongs to $\mathcal{R}$. Fix real numbers $\alpha, \beta \in [0,1]$. Then for N chosen such that*

$$\sum_{i=0}^{d-1} \binom{N}{i} \alpha^i (1-\alpha)^{N-i} \leq \beta$$

the optimal solution of the sampled stochastic program (4.5) *is feasible for* (4.4) *with probability no smaller than* $1-\beta$.

It is known that this result can be improved, e.g., by using sampling-and-discarding [32] and multi-stage or iterative schemes [28, 34, 38]. In this paper we stick with the requirement in Theorem 4.1 to illustrate our proposed strategies; improvements can be made accordingly by modifying the use of Theorem 4.1 to better results available in the literature.

4.3.2. Scenario Generation Under Unknown Distribution: A Bayesian Perspective

The scenario generation approach depicted in Theorem 4.1 requires direct observations on ξ or the capacity to obtain Monte Carlo samples for ξ. In problems with learning, the distribution of ξ is not fully known, and Theorem 4.1 does not apply directly. We shall adopt a Bayesian perspective that naturally integrates to many online algorithms (e.g., Thompson sampling). Suppose ξ follows a parametric distribution $G|\mu$ with unknown parameter μ. After specifying a prior distribution for μ and collecting some data historically, we have a posterior distribution for μ denoted by F. We seek to use Monte Carlo sampling to conduct an analog of scenario generation so that a posterior credibility guarantee

$$P_{\mu}(P_{\xi|\mu}(f(x,\xi)\in\mathcal{A})\geq 1-\alpha)\geq 1-\beta \tag{4.6}$$

is achieved, where P_{μ} denotes the posterior probability distribution on μ, and $P_{\xi|\mu}$ denotes the distribution of ξ given a parameter value of μ. In other words, we want the chance constraint to hold with a posterior credibility level $1-\beta$. Note that this is a natural Bayesian analog of the frequentist result in Theorem 4.1. In the setting of Theorem 4.1, P

is not known but data are available, so that a $1-\beta$ confidence is attained. In our current Bayesian investigation, P is not known but subject to a posterior belief summarized by the distribution of μ, and we want this posterior credibility to be $1-\beta$.

Directly sampling ξ using any particular value of μ does not sufficiently capture the posterior uncertainty. To blend the latter into a scenario generation, we can use a two-level sampling, where in the first level we generate a posterior sample for μ, and in the second level we generate ξ conditional on μ. This sampling procedure is described in Algorithm 1.

Algorithm 1 Posterior Constraint Sample Generator (PCSG)

1. Repeat N times:
   a) Generate $\mu^n \sim F$.
   b) Generate $\xi^n \sim G|\mu^n$.
2. Impose the constraints $f(x,\xi^n) \in \mathcal{A},\ n=1,\dots,N$ and solve the sampled program

$$\begin{aligned} \max_{x} \quad & h(x) \\ \text{subject to} \quad & f(x,\xi^n) \in \mathcal{A},\ \forall\ n=1,\dots,N \end{aligned} \tag{4.7}$$

In PCSG we encounter two different sources of randomness. First is the statistical noise from the uncertainty of μ, captured by the posterior credibility level $1-\beta$. The second source is the Monte Carlo error, and we denote by $1-\delta$ the confidence level induced from this error. By choosing a suitable sample size N in terms of α, β, δ, PCSG turns out to achieve a guarantee below.

Theorem 4.2. *Suppose $f(x,\xi) \in \mathcal{A}$ is a convex set in $x \in \mathbb{R}^d$, and $h(x)$ is concave. Suppose also that any instance of (4.7) belongs to $\mathcal{R}$. Fix real numbers $\delta, \alpha, \beta \in [0,1]$ and choose*

$$\sum_{i=0}^{d-1} \binom{N}{i} (\alpha\beta)^i (1-\alpha\beta)^{N-i} \leq \delta \tag{4.8}$$

Consider a solution x obtained from the sampled program in PCSG. Then,

1. *x satisfies* (4.6) *with a Monte Carlo confidence $1-\delta$, i.e.,*

$$P_{MC}(P_{\mu}(P_{\xi|\mu}(f(x,\xi)\in\mathcal{A})\geq 1-\alpha)\geq 1-\beta)\geq 1-\delta \tag{4.9}$$

   where the outermost P_{MC} denotes the probability with respect to the N Monte Carlo samples.

2. *x satisfies*

$$E_{MC}[P_{\mu}(P_{\xi|\mu}(f(x,\xi)\in\mathcal{A})\geq 1-\alpha)]\geq (1-\beta)(1-\delta) \tag{4.10}$$

   where $E_{MC}[\cdot]$ denotes the expectation with respect to the N Monte Carlo samples.

Theorem 4.2 Part 1 stipulates that choosing N in (4.8) achieves chance-constraint feasibility with a Bayesian credibility $1-\beta$, under a Monte Carlo confidence $1-\delta$. Part 2 will be useful in generalizing to the multi-stage setting presented next.

4.3.3. Sequential Policies

We now move our analysis to the sequential problem depicted in (4.3). We generalize PCSG, with the posterior update occurring at every step of the horizon and the sample size required at each step modified in order to achieve a chance constraint guarantee over the whole horizon. Let N_t be the sample size used in step t, which depends on α and also the confidence-level parameters β_t and δ_t. We denote F_0 as the prior distribution of μ and F_t as the posterior distribution of μ at step t. We denote $G_t|\mu$ as the distribution of ξ_t given μ. Note that μ is a parameter shared among the ξ_t at different steps so that information can be learned over time. To distinguish the real data from the Monte Carlo samples, we use $\tilde{\xi}_{1:(t-1)}$ to denote the actual data of ξ coming from steps 1 to $t-1$.

We have the following procedure:

Algorithm 2 Dynamic PCSG

Set F_0 as the prior distribution of μ. For $t = 1, \ldots, T$: While $t \in \mathcal{S}$, and given F_t and the realized $x_{1:(t-1)}$ and $\tilde{\xi}_{1:(t-1)}$:

1. Repeat N_t times:
   a) Generate $\mu^n \sim F_t$.
   b) Generate $\xi^n \sim G_t|\mu^n$.
2. Impose the constraints

$$f_t(x_{1:(t-1)}, x_t, \tilde{\xi}_{1:(t-1)}, \xi^n) \in \mathcal{A}_t, \ \forall\ n = 1, \ldots, N_t \tag{4.11}$$

   at stage t.

It is understood that in the second step of Dynamic PCSG, the constraints are imposed together with an appropriate objective function (typically the cost-to-go in formulation (4.3)) to form a stepwise optimization with decision variable x_t. The following result gives the choice of N_t and the resulting guarantee:

Theorem 4.3. *Suppose the stepwise safety conditions are all convex sets, the objective function at every step is concave, and $x_t \in \mathbb{R}^d$. Suppose also that any instance of the optimization resulted from imposing* (4.11) *belongs to $\mathcal{R}$. Suppose $0 \leq \beta_t, \delta_t \leq 1$ are constants such that*

$$\sum_{i=0}^{d-1} \binom{N_t}{i} (\alpha\beta_t)^i (1 - \alpha\beta_t)^{N_t - i} \leq \delta_t \tag{4.12}$$

and

$$\sum_{t \in \mathcal{S}} (\beta_t + \delta_t - \beta_t\delta_t) \leq \beta\lambda \tag{4.13}$$

Then the policy obtained from Dynamic PCSG is feasible for (4.3) *under the updated pos-*

terior distribution with probability at least $1-\lambda$*, i.e.,*

$$P_{MC}(P_{\mu_{1:T}}(P_{\xi_t|\mu_t}(f_t(x_t,\xi_t)\in\mathcal{A}_t|\mathcal{F}_{t-1})\geq 1-\alpha,\forall t\in\mathcal{S})\geq 1-\beta)\geq 1-\lambda \tag{4.14}$$

where $P_{\mu_{1:T}}$ *denotes the probability with respect to* $\mu_1,\ldots,\mu_t$*, where each* $\mu_t\sim F_t$*, the posterior distribution of* μ *at step* t*, and* $P_{\xi_t|\mu_t}$ *denotes the probability with respect to* ξ_t *given a realized parameter of* μ_t*.*

Theorem 4.3 asserts that the round-specific statistical parameters, namely the posterior credibility $1-\beta_t$ and the Monte Carlo confidence level $1-\delta_t$, which determine the constraint sample size, can be chosen to satisfy a local condition (4.12) and a global condition (4.13) to achieve an overall statistical guarantee.

For convenience we can set $\beta_t=\delta_t$ and both equal to some constant, say γ_t. This γ_t can be set to be stage-independent or dependent. The following subsection shows two choices of γ_t.

4.3.4. Two Explicit Strategies

We demonstrate two choices of $\{N_t\}$ in terms of $\{\gamma_t\}$. The first choice is a simple one that requires knowledge of the horizon length T, by setting γ_t to be a constant. The second choice uses a decaying γ_t, consequently an increasing sample size N_t, which does not require knowledge of T a priori. For the first strategy, we have:

Proposition 4.1. *Given a time horizon* T*, if we let* $\beta_t=\delta_t=\gamma$ *for all* $t\in\mathcal{S}$ *such that* $\gamma\leq 1-\sqrt{1-\beta\lambda/|\mathcal{S}|}$*, then* (4.14) *holds.*

The following describes our second strategy that is stage-dependent such that (4.14) holds without knowing the horizon T or $|\mathcal{S}|$ a priori:

Proposition 4.2. *If we let $\beta_t = \delta_t = \gamma_t$ for all $t \in \mathcal{S}$ such that $\gamma_t = (1/\zeta(t)^\rho) \wedge \eta$, where $\rho > 1$, $0 < \eta < 1$, and $\zeta(t) = \sharp\{s \in \mathcal{S} : s \leq t\}$ (i.e., $\zeta(t)$ is the "counter" of t in $\mathcal{S}$) such that*

$$2\eta^{1-1/\rho} + \frac{2}{\rho - 1}\frac{1}{(1/\eta^{1/\rho} - 1)^{\rho-1}} - \frac{\eta^2}{\eta^{1/\rho} + 1} - \frac{1}{2\rho - 1}\frac{1}{(1/\eta^{1/\rho} + 1)^{2\rho-1}} \leq \beta\lambda \qquad (4.15)$$

then (4.14) *holds regardless of* $|\mathcal{S}|$.

For example, if ρ is set to be 2, then (4.15) becomes $2\sqrt{\eta} + 2\sqrt{\eta}/(1 - \sqrt{\eta}) - \eta^2/(\sqrt{\eta} + 1) - \eta^{3/2}/(3(1 + \sqrt{\eta})^3) \leq \beta\lambda$.

4.4. Integration Into Thompson Sampling

We illustrate the integration of our strategies with Thompson sampling, which also operates via Bayesian updating, by an example of revenue maximization in online advertising [103]. The advertiser is interested in maximizing the expected revenue across a portfolios of keywords or biddable ad units while ensuring that the budget constraint is not violated. When the advertiser selects a bid value for a keyword it results in ad clicks, the volume of which is stochastic. The distribution of clicks and the associated revenue is not initially known to the decision-maker and needs to be learned over time. Further, the cost associated with the choice of a bid is also unknown and hence, there is uncertainty regarding how the budget will be affected.

To be more concrete, consider a set of K bid values $\{\kappa_1, ..., \kappa_K\}$, for M items labeled $\{\pi_1, ..., \pi_M\}$, over the campaign horizon T. Bidding value j on item i will induce an average revenue r_{ij} and cost c_{ij} respectively. These quantities are assumed to follow some independent Poisson distributions with initially unknown parameters (the Poisson assumptions come from the click count nature). In each period $t = 1, \ldots, T$, the advertiser picks a bid

value j from every item, observes the outcome, i.e., the realizations of $r_{ij}, c_{ij}, i = 1, \ldots, M$. She gains $\sum_{i=1}^{M} \sum_{j=1}^{K} r_{ij} x_{ij}$ and consumes $\sum_{i=1}^{M} \sum_{j=1}^{K} c_{ij} x_{ij}$ from the budget, where x_{ij} is the allocation portion for bid value j of item i (i.e., the fraction of time or the probability in a randomized scheme that is allocated to this particular bid value and item).

The advertiser's goal is to maximize the total revenue while maintain a budget constraint with high probability. In other words, letting $r_{ij}(t)$ and $c_{ij}(t)$ be the realized revenue and cost for a bid at time t, and $x_{ij}(t)$ be the corresponding allocation variables, she wants to maximize $E[\sum_{t=1}^{T} \sum_{i=1}^{M} \sum_{j=1}^{K} r_{ij}(t) x_{ij}(t)]$. A typical budget constraint is a bound given by the remaining budget averaged over the remaining horizon. Specifically, let the overall budget be B. Denoting $B(t-1) = B - \sum_{u=1}^{t-1} \sum_{i=1}^{M} \sum_{j=1}^{K} c_{ij}(u) x_{ij}(u)$ as the remaining budget before epoch $t \in \mathcal{S}$, the advertiser wants to keep $P(\sum_{u=1}^{t} \sum_{i=1}^{M} \sum_{j=1}^{K} c_{ij}(u) x_{ij}(u) \leq B(t-1)/(T-t+1)) \geq 1-\alpha \; \forall t \in \mathcal{S}$. In the following, we will concentrate on this particular budget feasibility requirement.

4.4.1. A Budgeted Algorithm

The setting of this problem resembles a recent study [58] on a network revenue management problem. They developed a Thompson sampling algorithm to sequentially assign a price vector to items under resource constraints, where each step involves a knapsack optimization problem. Here we present a variant of their algorithm to suit our setting (Algorithm 3). This initial algorithm does not take into account the possibility of constraint violation; the idea is to later illustrate how our Dynamic PCSG strategy can be integrated.

Denote by $X_{ij}(t-1)$ the allocation units on bid value j for item i cumulated in the first $t-1$ rounds, and denote by $W_{ij}^{r}(t-1)$ and $W_{ij}^{c}(t-1)$ the total revenue and cost generated by assigning bid value j to item i during these periods respectively. In Algorithm 3, the advertiser samples from the joint posterior distributions of θ_{ij}, the unknown Poisson rate

of the revenue, and μ_{ij}, the rate of the cost, corresponding to bid value j for item i. We put independent Gamma prior distributions for these parameters and hence the posterior distributions are also independent. The posterior samples of these parameters are then used in a linear program to decide the allocation.

Algorithm 3 Budgeted Thompson Sampling for deterministically constrained problems adopted from [58]

Given a total budget $B(0) = B$. For $t = 1, ..., T$, do the following:

1: For each bid value j and each item i, sample θ_{ij} from $Gamma(W^r_{ij}(t-1)+1, X_{ij}(t-1)+1)$ and μ_{ij} from $Gamma(W^c_{ij}(t-1)+1, X_{ij}(t-1)+1)$.

2: Solve the following linear program:

$$\begin{aligned} \max_{x} \quad & \sum_{j=1}^{K}\sum_{i=1}^{M} \theta_{ij} x_{ij} \\ \text{subject to} \quad & \sum_{j=1}^{K}\sum_{i=1}^{M} \mu_{ij} x_{ij} \leq \frac{B(t-1)}{T-t+1} \\ & \sum_{j=1}^{K} x_{ij} \leq 1, \ \forall\, i = 1, \ldots, M \\ & x_{ij} \geq 0, \ \forall\, i = 1, \ldots, M, \ j = 1, \ldots, K \end{aligned}$$

to obtain $(x^*_{ij}(t))_{i=1,...,M, j=1,...,K}$.

3: The revenue, $r_{ij}(t)$, and the cost, $c_{ij}(t)$, generated by assigning bid value j on item i are revealed. We update $X_{ij}(t) = X_{ij}(t-1) + x^*_{ij}(t)$, $W^r_{ij}(t) = W_{ij}(t-1) + r_{ij}(t)x^*_{ij}(t)$, $W^c_{ij}(t) = W^c_{ij}(t-1) + c_{ij}(t)x^*_{ij}(t)$ and $B(t) = B(t-1) - \sum_{j=1}^{K}\sum_{i=1}^{M} c_{ij}(t)x^*_{ij}(t)$.

4: If $B(t) \leq 0$, the algorithm terminates.

4.4.2. Chance Constrained Budgeted Thompson Sampling

We want to ensure $P(\sum_{u=1}^{t}\sum_{i=1}^{M}\sum_{j=1}^{K} c_{ij}(u)x_{ij}(u) \leq B(t-1)/(T-t+1)) \geq 1-\alpha \ \forall t \in \mathcal{S}$ holds with posterior credibility $1-\lambda$. To achieve this, we integrate our Dynamic PCSG into Step 2 of Algorithm 3 by restructuring the involved optimization program. Algorithm 4 shows Dynamic PCSG in this particular setting.

The output of this procedure is a set of constraints, which will be used in the linear

Algorithm 4 Dynamic PCSG for the Bidding Problem

1. Set $F_{ij}(0), i = 1, \ldots, M, j = 1, \ldots, K$ as the prior distribution of μ_{ij}. For $t = 1, \ldots, T$: Given $B(t-1)$ and $F_{ij}(t-1)$, the posterior distribution of μ_{ij} given $\mathcal{F}_{t-1}$,
   a) Repeat N_t times:
      i. Generate $\mu_{ij}^n \sim F_{ij}(t-1)$ for each item $i = 1, \ldots, M$ and bid value $j = 1, \ldots, K$.
      ii. Generate $\xi_{ij}^n \sim Poisson(\mu_{ij}^n)$ for each item $i = 1, \ldots, m$ and bid value $j = 1, \ldots, K$.
   b) Form the constraints

$$\sum_{j=1}^{K} \sum_{i=1}^{M} \xi_{ij}^n x_{ij} \leq \frac{B(t-1)}{T-t+1}, \ \forall \ n = 1, \ldots, N_t$$

program of the budgeted Thompson sampling. Algorithm 5 shows how Dynamic PCSG can be integrated into Algorithm 3. The following is an immediate consequence of Theorem 4.3:

Corollary 4.4. *Suppose that any instance of the sampled program in* (4.16) *belongs to* $\mathcal{R}$. *Suppose* $0 \leq \beta_t, \delta_t \leq 1$ *are chosen to satisfy*

$$\sum_{i=0}^{KM-1} \binom{N_t}{i} (\alpha\beta_t)^i (1-\alpha\beta_t)^{N_t-i} \leq \delta_t$$

and

$$\sum_{t \in \mathcal{S}} (\beta_t + \delta_t - \beta_t\delta_t) \leq \beta\lambda$$

Consider a modification of Algorithm 5 such that at any point of time, if the total budget B *is fully depicted, we refill the shortfall and add extra budget* B *(so that the total remaining budget in the next step returns to the full level* B*). Then the sequence of decisions obtained*

Algorithm 5 Chance-constrained Thompson Sampling (CCTS)

Initialize $\alpha, \beta_t, \delta_t \in [0,1]$ satisfying (4.13). Given a total budget $B(0) = B$. For $t = 1, ..., T$, do the following:

1: For each bid value j and each item i, sample θ_{ij} from $Gamma(W_{ij}^r(t-1)+1, X_{ij}(t-1)+1)$ and μ_{ij} from $Gamma(W_{ij}^c(t-1)+1, X_{ij}(t-1)+1)$.

2: Run Dynamic PCSG using N_t samples to get the constraints

$$\sum_{j=1}^{K}\sum_{i=1}^{M} \xi_{ij}^n x_{ij} \leq \frac{B(t-1)}{T-t+1},\ \forall\ n = 1, \ldots, N_t$$

3: Solve the following linear program:

$$\begin{aligned} \max_{x} \quad & \sum_{j=1}^{K}\sum_{i=1}^{M} \theta_{ij} x_{ij} \\ \text{subject to} \quad & \sum_{j=1}^{K}\sum_{i=1}^{M} \xi_{ij}^n x_{ij} \leq \frac{B(t-1)}{T-t+1},\ \forall\ n = 1, \ldots, N_t \\ & \sum_{j=1}^{K} x_{ij} \leq 1,\ \forall\ i = 1, \ldots, M \\ & x_{ij} \geq 0,\ \forall\ i = 1, \ldots, M,\ j = 1, \ldots, K \end{aligned} \tag{4.16}$$

to obtain $(x_{ij}^*(t))_{i=1,...,M, j=1,...,K}$.

4: The revenue, $r_{ij}(t)$, and the cost, $c_{ij}(t)$, generated by assigning bid value j on item i are revealed. We update $X_{ij}(t) = X_{ij}(t-1) + x_{ij}^*(t)$, $W_{ij}^r(t) = W_{ij}(t-1) + r_{ij}(t)x_{ij}^*(t)$, $W_{ij}^c(t) = W_{ij}^c(t-1) + c_{ij}(t)x_{ij}^*(t)$ and $B(t) = B(t-1) - \sum_{j=1}^{K}\sum_{i=1}^{M} c_{ij}(t)x_{ij}^*(t)$.

5: If $B(t) \leq 0$, the algorithm terminates.

will satisfy

$$P_{\boldsymbol{\mu}_t}\left(P_{\mathbf{c}_t|\boldsymbol{\mu}_t}\left(\sum_{u=1}^{t}\sum_{i=1}^{M}\sum_{j=1}^{K} c_{ij}(u)x_{ij}(u) \leq \frac{B(t-1)}{T-t+1}\right) \geq 1-\alpha\ \forall t \in \mathcal{S}\right) \geq 1-\beta$$

with Monte Carlo confidence level at least $1-\lambda$*, where* $P_{\boldsymbol{\mu}_t}$ *denotes the probability with respect to all* μ_{ij} *under the posterior distributions* $F_{ij}(t-1)$*, and* $P_{\mathbf{c}_t|\boldsymbol{\mu}_t}$ *denotes the probability with respect to all* c_{ij} *given the realizations of* μ_{ij}*'s.*

We mention that the "modification of Algorithm 5" introduced in Corollary 4.4 is only

a technicality that takes care of the unusual situation when the entire available budget is prematurely depleted. Since we divide the remaining budget by the remaining horizon (a common practice to set per-round budgets) to form our constraint at each step, the scenario of total budget depletion before the last step rarely happens.

4.5. Numerical Results

We examine the empirical performance of our proposed strategy on a synthetic dataset with two items and three bid values ($M = 2$, $K = 3$) over the time horizon $T = 100$. The cost and revenue of each item-bid value pair follow Poisson distributions with parameters taken uniformly from an interval that is calibrated from a real data set owned by Adobe Systems, Inc. We test with five different values of the overall budget $B = (a \sum_{i=1}^{M} \bar{\lambda}_i) \times T$ where $a \in [0.5, 0.75, 1, 1.25, 1.5]$ and $\bar{\lambda}_i$ is the average cost of item i over the K bid values. This choice of B roughly matches the scale of the total cost over the time horizon. The per-round budget is defined as the remainder of the overall budget divided by the remaining number of rounds.

To test our chance-constrained Thompson sampling (CCTS) in Algorithm 5, we use three different settings for $\mathcal{S}$, i.e. $\mathcal{S}_1 = \{25, 50, 75\}$, $\mathcal{S}_2 = \{20, 40, 60, 80, 100\}$ and $\mathcal{S}_3 = \{10, 20, 30, 40, 50, 60, 70, 80, 90, 100\}$. We also enforce $\alpha = 0.1$, $\beta = \lambda = 0.3$, and $\beta_t = \delta_t = \gamma$ where γ is taken as the upper bound depicted in Proposition 4.1. We conduct 500 simulation runs. For each run, we estimate the proportion of violation of the decision using 100 inner repetitions of ξ_t, at each step $t \in \mathcal{S}$. Figure 4.1 depicts the box-plots showing the distribution of the proportion of violation. For all the tested budget levels and choices of $\mathcal{S}$, CCTS was able to maintain the proportion of budget violation well below the 10% tolerance at the relevant steps. This implies, moreover, that the overall violation (i.e., at least one violation at a step in $\mathcal{S}$) is also below 10%.

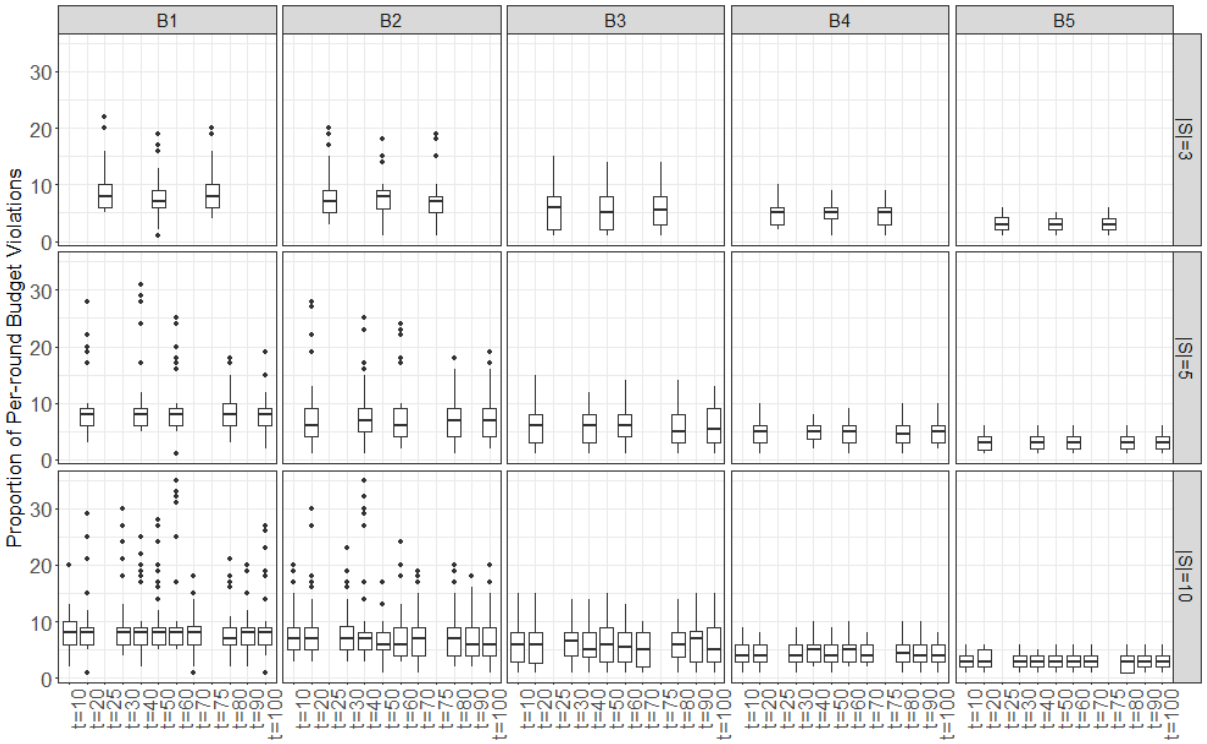

Figure 4.1: Estimated proportion of violation at each step in $\mathcal{S}$, for $\mathcal{S} = \mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3$ and five different budget levels, using 500 simulation runs under CCTS.

We also test the performance of CCTS against the following algorithms: 1) a hypothetical algorithm that assumes the distribution of the costs and revenues are all known and draws Monte Carlo samples from it, otherwise the same as CCTS; 2) the deterministically constrained Thompson sampling (DCTS) in Algorithm 3; 3) the algorithm in [9] that uses reward-to-cost ratios; and 4) [21] that uses an initial learning phase (in our experiment we set the learning phase to 50 steps). Figure 4.2 shows the distributions of the proportion of violations at $t \in \mathcal{S}_2$ based on 500 simulation runs, for the five described budget levels. We see that only CCTS and the hypothetical algorithm maintains the proportion of violation to well below 10%, whereas the other three algorithms fluctuate around 20-40%, as they do not account for the chance constraint on the per-round budget. On the other hand, Figure 4.3 shows the average cumulative revenue achieved through the horizon (the bar

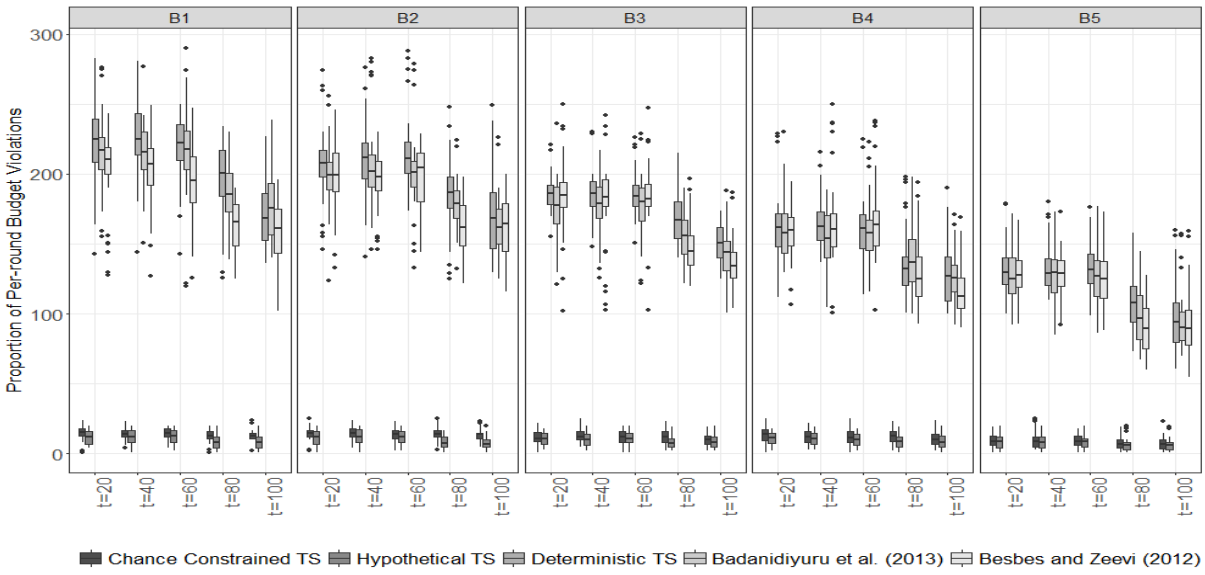

Figure 4.2: Estimated proportion of violation at each step in $\mathcal{S}_2$, with 500 simulation runs for different algorithms

depicts one standard deviation around the average). DCTS appears the best in terms of cumulative rewards, and [9] and [21] perform similarly. CCTS achieves less rewards (around 15%), which can be viewed as the price of maintaining the chance constraint. The hypothetical algorithm performs better than CCTS, not surprisingly given the full distributional knowledge. This behavior persists for $\mathcal{S}_1$ and $\mathcal{S}_3$ as well as for several priors we have tested (similar to plots Figures 4.2 and 4.3).

The above experiments all used the budget violations calculated using the true underlying distribution of the cost. In Table 4.1, we compare with the calculation based on the evolving posterior distributions, in terms of the average of all the proportions of violations for $t \in \mathcal{S}_1, \mathcal{S}_2$ and $\mathcal{S}_3$ at the five described budget levels. The average proportion of violation based on the posterior distribution is consistently lower than the one based on the true distribution for all budget levels and $\mathcal{S}$. This is expected since our chance constraint is maintained under the posterior distribution (as Theorem 4.3 states). However, the proportion of violations is maintained below 10% under the true distribution, thanks to the

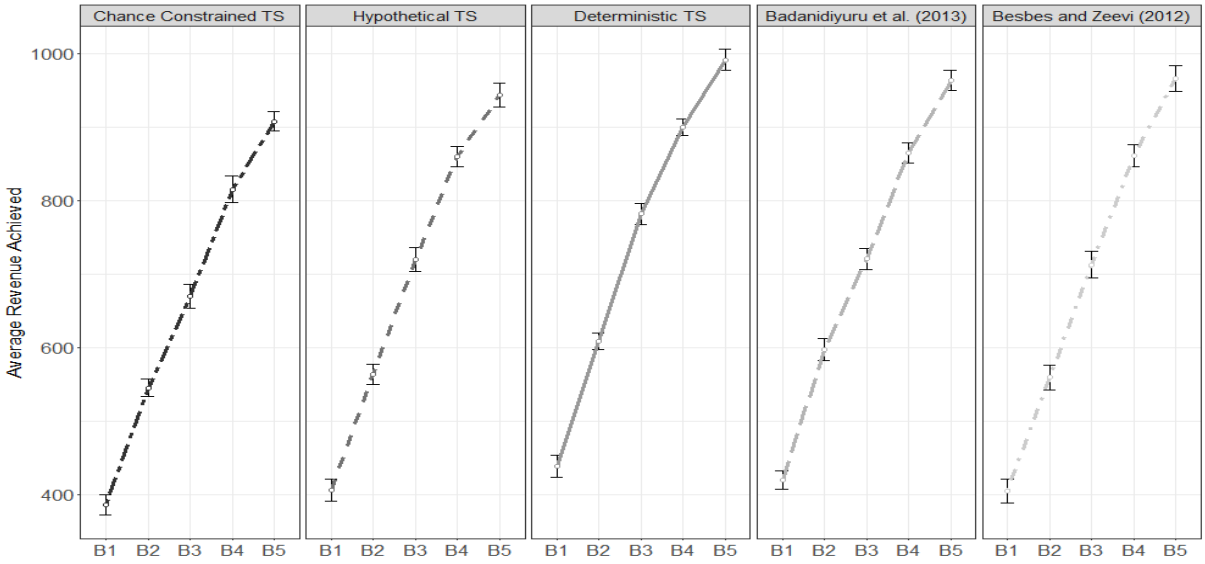

Figure 4.3: Average cumulative revenue over $T = 100$ achieved by different algorithms under $\mathcal{S}_2$

relatively conservative performance of our approach in abiding with the chance constraint.

Table 4.1: Comparison of the average proportions of violations based on the true distribution and the updated posterior distributions over 500 simulation runs

Budget		$\mathcal{S}_1$	$\mathcal{S}_2$	$\mathcal{S}_3$
$B1$	True Dist.	0.018	0.019	0.018
	Posterior Dist.	0.011	0.012	0.009
$B2$	True Dist.	0.016	0.015	0.014
	Posterior Dist.	0.01	0.01	0.008
$B3$	True Dist.	0.014	0.013	0.013
	Posterior Dist.	0.008	0.009	0.007
$B4$	True Dist.	0.013	0.011	0.01
	Posterior Dist.	0.009	0.007	0.005
$B5$	True Dist.	0.008	0.008	0.007
	Posterior Dist.	0.006	0.005	0.004

In addition to the number of violated budget constraints, we investigated the amount of budget violation at $t \in \mathcal{S}$. Figure 4.4 depicts the distributions of the total amounts

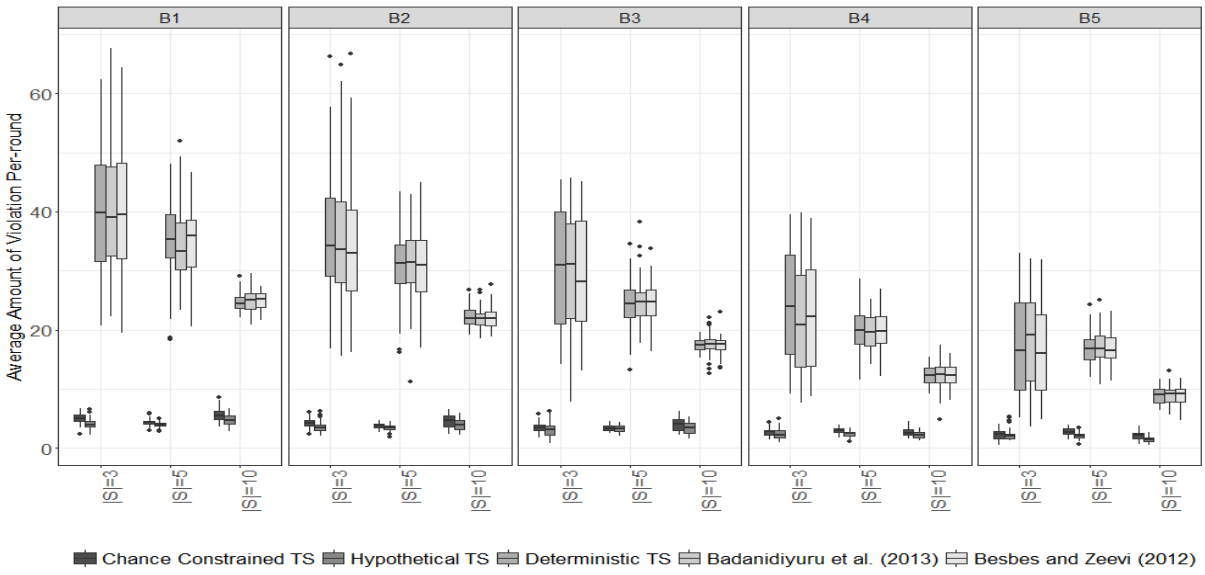

Figure 4.4: Total Amount of budget violations occurred over $\mathcal{S}_1, \mathcal{S}_2$ and $\mathcal{S}_3$ for different algorithms

of budget violation in $\mathcal{S}_1, \mathcal{S}_2$ and $\mathcal{S}_3$ for the five budget levels over 500 simulation runs. Similar to the proportion of budget constraint violations, the average amounts of violation for CCTS and DTS tend to be much lower than the other three algorithms, with at most 25% of those of the other three algorithms in the same setting. This suggests a strong dependence between the proportion and the amount of violation which substantiates the use of CCTS in maintaining over-spending even in the monetary scale.

4.6. Conclusion

We have studied sequential learning subject to constraints that need to be satisfied with high probability. We investigated a methodology to obtain posterior statistical guarantees for the feasibility of these constraints, by generalizing the constraint sampling approach in chance-constraint programming to a two-level Monte Carlo procedure and analyzing the sample size needed in achieving overall feasibility through the learning horizon. We

further incorporated our scheme into Thompson sampling using an online advertisement example. As far as we know, this work represents the first methodological investigation of "soft" stochastic constraints in sequential learning. In subsequent work, we will investigate the tightening of the requirements in sampled constraints, via for instance analyzing the correlation among decisions at different stages, and will also study the scalability of this approach to higher-dimensional problems.

Chapter 5.

Conclusions and Future Research

In this dissertation we laid out several frameworks for integrated predictive modeling and optimal decision making processes in the context of operations management. In the main three chapters we studied different designs for reliable abstraction of the overall available data into personalized operational decision making process. From separated estimation and optimization toward individualized optimization under uncertainty and sequential decision making under probabilistic constraints we were able to establish practical tools with theoretical and empirical performance guarantees.

In the second chapter, we introduced a new perspective to address the admission policy for a network of critical care units. We took advantage of one of the early measures of patient mortality risk to show how patient risk metrics can be used to rationalize and optimize the admission of patients given the limited capacity of higher level of care in a hospital. We proposed a methodology that generates optimal admission policies for different ICUs and IMCs in a hospital. Considering the significant cost of critical care units, our model can guide hospital admission policies towards a data-informed solution that guarantees a higher service level as well as higher accessibility. In fact, the numerical results based on real data from St. Joseph Mercy Hospital of Ann Arbor proved significant

improvement in the overall hospital's service rate and better patient placement. Our proposed methodology is easily generalizable and can be adapted by any hospital with a real-time health risk metric and a suitable patient records database. We expect this approach to gain acceptance as hospitals deploy methods to assess patients' risks. Our model was shown to diminish wait times to admit high-risk patients to higher level of care by a more appropriate initial admission decision upon patient's arrival. Hence, it can potentially decrease mortalities caused by delayed access or non-access to critical care units.

For future work, one direction is to consider dynamically updated patient mortality/health risk assessments, provided such longitudinal data are available. Dynamic risk measures can provide significant information regarding the patient recovery process that is valuable for improved admission decisions. Such data supports the medical decision making processes by prescribing meaningful interventions throughout the patient trajectory in different layers of the care delivery system. Another interesting direction is to link the admission decisions to the readmission rate of the hospital. One can design a similar methodology in which the readmission risk of the patient is the main contributor to the admission decision. In such a case, post-discharge interventions can also be tailored to the care delivery planning network so that a more comprehensive system design is achievable.

In the third chapter of this dissertation, we developed comprehensive individualized statistical models to predict surgical case duration, in which we considered multiple data sources, categories and interactions. These include information about patients, surgeons, and operational data. Advanced predictive modeling, including Natural Language Processing, was applied to gain insights from both structured and unstructured data. Utilizing a test set for validation of our proposed models at Memorial Sloan Kettering Cancer Center, we demonstrated superior predictive ability. The inclusion of textual data contributed to a

significant reduction in average estimation error. Our proposed statistical models are capable of generating reliable prediction intervals, serving as a new source of information that facilitates improved management of operating room resources, scheduling, and sequencing of surgeries. Finally, using the developed framework to achieve improved, more accurate prediction of case duration, we demonstrated generalizability across all surgical services at the institution.

Motivated by the initial predictive study of the surgical case duration, in the second part of the third chapter, we studied the fundamental appointment scheduling problem under job duration uncertainty in a prescriptive data-informed framework. Such problems are of particular interest in healthcare applications, e.g. surgery scheduling, where considerable amount of patient, staff and operational data is readily available prior to the surgery. Borrowing ideas from machine learning theory, we developed individualized modeling frameworks for the stochastic, robust and distributionally robust optimization settings and we demonstrated the statistical consistency of our proposed methodologies regarding the optimality guarantees. The implementation of our proposed individualized models on real data from MSKCC has shown impressive reduction in waiting and overtime costs compared to the traditional scheduling models.

Considering the individualization framework, there are many promising research paths one can investigate. The boost of data, specially in service and manufacturing sectors, has paved the way for further theoretical and empirical studies to incorporate machine learning and statistical tools in operations management. One future perspective can be related to a more in-depth analysis of ensemble learning algorithms. In Chapter 3, we elaborated the power of bagging methods and their flexibility in integration with optimal decision making. However, inferring the statistical characteristics (variance, confidence intervals, etc.) of ensemble learning models such as random forests is a challenging task

and their application in simulation-based modeling can be tricky. In a recent paper, we considered simulation-based estimation where the inputs are calibrated from predictive models generated by random forests [97]. We investigated the estimation of the output variance contributed from the noises in the input prediction, which can be used to construct output confidence intervals. In particular, we studied the integration of simulation runs with infinitesimal jackknife estimator that can reduce the computational burden from double layers of bootstrapping. Nevertheless, the integration of such inferential methods and optimal decision making is yet to be investigated.

Finally in the fourth chapter, we focused on sequential learning subject to constraints that need to be satisfied with high probability. We investigated a methodology to obtain posterior statistical guarantees for the feasibility of these constraints, by generalizing the constraint sampling approach in chance-constraint programming to a two-level Monte Carlo procedure and analyzing the sample size needed in achieving overall feasibility through the learning horizon. We further incorporated our scheme into Thompson sampling in the context of digital marketing campaigns. To the best of our knowledge, this work represents the first methodological investigation of "soft" stochastic constraints in sequential learning.

In subsequent work, based on our sequential learning under probabilistic constraints platform, we aim to investigate the tightening of the requirements in sampled constraints, via analyzing the correlation among decisions at different stages. We also plan to study the scalability of this approach to higher-dimensional problems. The novel framework presented in Chapter 4 provides enormous opportunity for further explorations in this area. For instance, once can envision a new rendition of the problem by introducing a penalized objective function that discourages over-spending in a sequential fashion. This may lead to concrete analytical results on bounding the overall expected amount of over-spending throughout the horizon which naturally complements our analysis on bounding the number

of violated rounds. Additionally, the integration of ideas discussed in Chapter 3 and 4 can also lead to promising findings. Incorporating the concept of quantiles and quantile regression, one can design risk-averse sequential decision making models by investigating the tail of revenue/cost distribution. These models are particularly desirable for portfolio optimization and medical decision making where the decision maker has to account for performance uncertainty in real time.

The results presented in this dissertation are a direct outcome of close engagement with several medical collaborators from cutting-edge healthcare organizations (i.e., University of Michigan Health System, St. Joseph Mercy Hospital and the Memorial Sloan Kettering Cancer Center) and our industry collaborators form Adobe Systems, Inc. These collaborations enabled us to identify interesting and challenging real-world problems, obtain relevant datasets, and ensure practical significance. The ideas and methods proposed in this dissertation, through comprehensive inclusion of data and elevated personalization, help reduce the cost of operations by establishing frameworks to achieve higher accessibility and service level, robust and efficient performance under uncertainty and reliable sequential decision making.

Appendix A.

Supplements to Chapter 2

A.1. Notations and MIP Models

Variables:

$\theta^{\mathbf{k}}_{\mathbf{udi}}$ binary variable that is equal to 1 if we allow patients of class k and risk group i to be admitted to unit u on day d, and equal to 0 otherwise (*decision variable*). $\mathbf{\Theta}$ refers to the matrix $[\theta^{\mathbf{k}}_{\mathbf{udi}}]$.

$\mu_{\mathbf{du}}$ mean census in unit u on day d, given $\mathbf{\Theta}$.

$\mathbf{v_{du}}$ variance of the census in unit u on day d, given $\mathbf{\Theta}$.

Parameters:

C_{du}	capacity of unit u on day d.
$\bar{F}_{dki}(a)$	probability that there are more than a arrivals of class k and risk group i on day d.
λ_{di}^{k}	mean arrival rate of patients of class k and risk group i on day d.
$\beta_{\mathbf{udi}}^{\mathbf{k}}$	maximum allowed number of admissions for patients of a certain class with a certain risk level to each unit on each day.
$P_{dkui}^{u'}(r)$	probability that a patient of class k and risk group i who arrived on day d is in unit u r days after admission, given that they were initially admitted to unit u'. By convention, $P_{dkui}^{u'}(r) = 0$ if $r < 0$.
p_j	probability that demand falls in the j^{th} demand interval.
R_i	expected risk score for a patient in risk group i.
$\hat{\sigma}_{du}$	observed standard deviation of the census in unit u on day d.
z_α	the coefficient associated with service level α.

Sets:

$\mathcal{D}, \mathcal{I}$	set of days of the week and set of risk groups, respectively.
$\mathcal{M}$	set of ICU-mandatory patients.
$\mathcal{M}'$	set of IMC-mandatory patients.
$\mathcal{K}$	set of patient classes. In this context, a patient's class is determined by her treatment needs and diagnosis. We will use patient class to determine which units a patient may be admitted to.
$\mathcal{U}$	set of units / wards.
$\mathcal{U}^{ICU}$	set of ICUs ($\mathcal{U}^{ICU} \subset \mathcal{U}$).
$\mathcal{U}^{IMC}$	set of IMCs ($\mathcal{U}^{IMC} \subset \mathcal{U}$).

The mathematical models are given below.

MIP model

$$\max_{\theta} \sum_{d\in\mathcal{D}} \sum_{k\in\mathcal{K}} \sum_{i\in\mathcal{I}\setminus\{\mathcal{M},\mathcal{M}'\}} \sum_{u\in\mathcal{U}\setminus\{\mathrm{GMB}\}} \theta^{\mathbf{k}}_{\mathbf{udi}} R_i \tag{A.1}$$

$$s.t. \quad C_{du} - \mu_{\mathbf{du}} \geq 0 \qquad \forall\ d\in\mathcal{D}, u\in\mathcal{U} \tag{A.2}$$

$$z^2_\alpha \mathbf{v_{du}} \leq C^2_{du} + \mu^{\mathbf{2}}_{\mathbf{du}} - 2C_{du}\mu_{\mathbf{du}} \qquad \forall\ d\in\mathcal{D}, u\in\mathcal{U} \tag{A.3}$$

$$\mu_{\mathbf{du}} = \sum_{i\in\mathcal{I}} \sum_{d'\in\mathcal{D}} \sum_{u'\in\mathcal{U}} \sum_{k\in\mathcal{K}} \sum_{n=0}^{\infty} \sum_{a=0}^{\infty} P^{u'}_{d'kui}(d-d'+7n) .\bar{F}_{d'ki}(a)\theta^{\mathbf{k}}_{\mathbf{u'd'i}} \qquad \forall\ d\in\mathcal{D}, u\in\mathcal{U} \tag{A.4}$$

$$\mu^{\mathbf{2}}_{\mathbf{du}} = \sum_{i\in\mathcal{I}} \sum_{d'\in\mathcal{D}} \sum_{u'\in\mathcal{U}} \sum_{k\in\mathcal{K}} \sum_{n=0}^{\infty} \sum_{a=0}^{\infty} \{P^{u'}_{d'kui}(d-d'+7n)\bar{F}_{d'ki}(a)\}^2 \theta^{\mathbf{k}}_{\mathbf{u'd'i}} + \sum_{i\in\mathcal{I}} \sum_{d'\in\mathcal{D}} \sum_{u'\in\mathcal{U}} \sum_{k\in\mathcal{K}} \sum_{\substack{i'\in\mathcal{I}\\ i'\neq i}} \sum_{\substack{d''\in\mathcal{D}\\ d''\neq d'}} \sum_{\substack{u''\in\mathcal{U}\\ u''\neq u'}} \sum_{\substack{k'\in\mathcal{K}\\ k'\neq k}} \theta^{\mathbf{k}}_{\mathbf{u'd'i}} \theta^{\mathbf{k'}}_{\mathbf{u''d''i'}} \qquad \forall\ d\in\mathcal{D}, u\in\mathcal{U} \tag{A.5}$$

$$\mathbf{v_{du}} = \sum_{i\in\mathcal{I}} \sum_{d'\in\mathcal{D}} \sum_{u'\in\mathcal{U}} \sum_{k\in\mathcal{K}} \sum_{n=0}^{\infty} \sum_{a=0}^{\infty} \{P^{u'}_{d'kui}(d-d'+7n)\bar{F}_{d'ki}(a) - [P^{u'}_{d'kui}(d-d'+7n)\bar{F}_{d'ki}(a)]^2\}\theta^{\mathbf{k}}_{\mathbf{u'd'i}} \qquad \forall\ d\in\mathcal{D}, u\in\mathcal{U} \tag{A.6}$$

$$\sum_{u \in \mathcal{U}} \theta^{\mathbf{k}}_{\mathbf{udi}} = 1 \qquad \forall\ d \in \mathcal{D}, i \in \mathcal{I}, k \in \mathcal{K} \tag{A.7}$$

$$\sum_{u \in \mathcal{U}^{ICU}} \theta^{\mathbf{k}}_{\mathbf{udi}} = 1 \qquad \forall\ d \in \mathcal{D}, i \in \mathcal{M}, k \in \mathcal{K} \tag{A.8}$$

$$\sum_{u \in \mathcal{U}^{IMC}} \theta^{\mathbf{k}}_{\mathbf{udi}} = 1 \qquad \forall\ d \in \mathcal{D}, i \in \mathcal{M}', k \in \mathcal{K} \tag{A.9}$$

$$\omega^{\mathbf{kk'}}_{\mathbf{u'u''d'd''ii'}} \leq \theta^{\mathbf{k}}_{\mathbf{u'd'i}} \qquad \forall\ d', d" \in \mathcal{D}, i, i' \in \mathcal{I}, u', u" \in \mathcal{U}, k, k' \in \mathcal{K} \tag{A.10}$$

$$\omega^{\mathbf{kk'}}_{\mathbf{u'u''d'd''ii'}} \leq \theta^{\mathbf{k'}}_{\mathbf{u''d''i'}} \qquad \forall\ d', d" \in \mathcal{D}, i, i' \in \mathcal{I}, u', u" \in \mathcal{U}, k, k' \in \mathcal{K} \tag{A.11}$$

$$\omega^{\mathbf{kk'}}_{\mathbf{u'u''d'd''ii'}} \geq \theta^{\mathbf{k}}_{\mathbf{u'd'i}} + \theta^{\mathbf{k'}}_{\mathbf{u''d''i'}} - 1 \qquad \forall\ d', d" \in \mathcal{D}, i, i' \in \mathcal{I}, u', u" \in \mathcal{U}, k, k' \in \mathcal{K} \tag{A.12}$$

$$\lambda^{k}_{di} \theta^{\mathbf{k}}_{\mathbf{udi}} \leq \beta^{\mathbf{k}}_{\mathbf{udi}} \qquad \forall\ d \in \mathcal{D}, i \in \mathcal{I}, u \in \mathcal{U}, k \in \mathcal{K} \tag{A.13}$$

$$\theta^{\mathbf{k}}_{\mathbf{udi}} \leq \theta^{\mathbf{k}}_{\mathbf{ud(i+1)}} \qquad \forall\ d \in \mathcal{D}, i \in \mathcal{I}, u \in \mathcal{U}^{ICU} \tag{A.14}$$

$$\theta^{\mathbf{k}}_{\mathbf{u'di}} \leq \theta^{\mathbf{k}}_{\mathbf{ud(i+1)}} \qquad \forall\ d \in \mathcal{D}, i \in \mathcal{I}, u' \in \mathcal{U}^{IMC}, u \in \mathcal{U}^{ICU} \tag{A.15}$$

$$\theta^{\mathbf{k}}_{\mathbf{u'di}} \leq \theta^{\mathbf{k}}_{\mathbf{ud(i+1)}} \qquad \forall\ d \in \mathcal{D}, i \in \mathcal{I}, u' \in \mathcal{U}^{GMB}, u \in \mathcal{U}^{IMC} \tag{A.16}$$

$$\mu_{\mathbf{du}}, \mathbf{v_{du}} \geq 0 \text{ and } \theta^{\mathbf{k}}_{\mathbf{udi}} \text{ binary} \qquad \forall\ d \in \mathcal{D}, i \in \mathcal{I}, u \in \mathcal{U} \tag{A.17}$$

MIP model with the Newton Approximation

$$\max_{\theta} \sum_{d\in\mathcal{D}} \sum_{k\in\mathcal{K}} \sum_{i\in\mathcal{I}\setminus\{\mathcal{M},\mathcal{M}'\}} \sum_{u\in\mathcal{U}\setminus\{\text{GMB}\}} \theta^{\mathbf{k}}_{\mathbf{udi}} R_i \tag{A.18}$$

$$s.t. \quad \mu_{\mathbf{du}} + \frac{z_\alpha}{2}\left(\frac{\mathbf{v_{du}}}{\hat{\sigma}_{du}} + \hat{\sigma}_{du}\right) \le C_{du} \qquad \forall\ d\in\mathcal{D}, u\in\mathcal{U} \tag{A.19}$$

$$\mu_{\mathbf{du}} = \sum_{i\in\mathcal{I}} \sum_{d'\in\mathcal{D}} \sum_{u'\in\mathcal{U}} \sum_{k\in\mathcal{K}} \sum_{n=0}^{\infty} \sum_{a=0}^{\infty} P^{u'}_{d'kui}(d-d'+7n)$$

$$.\bar{F}_{d'ki}(a)\theta^{\mathbf{k}}_{\mathbf{u'd'i}} \qquad \forall\ d\in\mathcal{D}, u\in\mathcal{U} \tag{A.20}$$

$$\mathbf{v_{du}} = \sum_{i\in\mathcal{I}} \sum_{d'\in\mathcal{D}} \sum_{u'\in\mathcal{U}} \sum_{k\in\mathcal{K}} \sum_{n=0}^{\infty} \sum_{a=0}^{\infty} \{P^{u'}_{d'kui}(d-d'+7n)\bar{F}_{d'ki}(a)$$

$$- [P^{u'}_{d'kui}(d-d'+7n)\bar{F}_{d'ki}(a)]^2\}\theta^{\mathbf{k}}_{\mathbf{u'd'i}} \qquad \forall\ d\in\mathcal{D}, u\in\mathcal{U} \tag{A.21}$$

$$\sum_{u\in\mathcal{U}} \theta^{\mathbf{k}}_{\mathbf{udi}} = 1 \qquad \forall\ d\in\mathcal{D}, i\in\mathcal{I}, k\in\mathcal{K} \tag{A.22}$$

$$\sum_{u\in\mathcal{U}^{ICU}} \theta^{\mathbf{k}}_{\mathbf{udi}} = 1 \qquad \forall\ d\in\mathcal{D}, i\in\mathcal{M}, k\in\mathcal{K} \tag{A.23}$$

$$\sum_{u\in\mathcal{U}^{IMC}} \theta^{\mathbf{k}}_{\mathbf{udi}} = 1 \qquad \forall\ d\in\mathcal{D}, i\in\mathcal{M}', k\in\mathcal{K} \tag{A.24}$$

$$\lambda^k_{di}\theta^{\mathbf{k}}_{\mathbf{udi}} \le \beta^{\mathbf{k}}_{\mathbf{udi}} \qquad \forall\ d\in\mathcal{D}, i\in\mathcal{I}, u\in\mathcal{U}, k\in\mathcal{K} \tag{A.25}$$

$$\theta^{\mathbf{k}}_{\mathbf{udi}} \le \theta^{\mathbf{k}}_{\mathbf{ud(i+1)}} \qquad \forall\ d\in\mathcal{D}, u\in\mathcal{U}^{ICU} \tag{A.26}$$

$$\theta^{\mathbf{k}}_{\mathbf{u'di}} \le \theta^{\mathbf{k}}_{\mathbf{ud(i+1)}} \qquad \forall\ d\in\mathcal{D}, u'\in\mathcal{U}^{IMC}, u\in\mathcal{U}^{ICU} \tag{A.27}$$

$$\theta^{\mathbf{k}}_{\mathbf{u'di}} \le \theta^{\mathbf{k}}_{\mathbf{ud(i+1)}} \qquad \forall\ d\in\mathcal{D}, u'\in\mathcal{U}^{GMB}, u\in\mathcal{U}^{IMC} \tag{A.28}$$

$$\theta^{\mathbf{k}}_{\mathbf{udi}} \text{ binary} \qquad \forall\ d\in\mathcal{D}, i\in\mathcal{I}, u\in\mathcal{U} \tag{A.29}$$

$$\mu_{\mathbf{du}}, \mathbf{v_{du}} \ge 0 \qquad \forall\ d\in\mathcal{D}, u\in\mathcal{U} \tag{A.30}$$

A.2. Proof of Theorems

Poof of Theorem 2.1 Let $L_{ski}^{au'}(r)$ be the independent and identically distributed (i.i.d.) stochastic location function of the a^{th} arrival of class k and risk group i on day s. The u' superscript denotes that this arrival was initially admitted to unit u'. $L_{ski}^{au'}(r) = (l_{ski}^{au'1}(r), l_{ski}^{au'2}(r), \ldots, l_{ski}^{au'N}(r))^T$, where $l_{ski}^{au'u}(r)$ equals 1 if the patient is in unit u r days after admission and equals 0 otherwise, with $\sum_{u=1}^{N} l_{ski}^{au'u}(r) \leq 1$, where N is the number of units. By convention, $l_{ski}^{au'u}(r) = 0$ if $r < 0$. Further, let A_{ks}^{i} be the stochastic number of patients of class k and risk group i who arrive on day s. Finally, let $B_{du}(t)$ be the stochastic census function of unit u on day d in week t, and let $\mathbf{e}_{\text{u}}$ be a row vector of zeros with a 1 in the u^{th} element such that

$$
\begin{aligned}
B_{du}(t) &= \sum_{i\in\mathcal{I}} \sum_{d'\in\mathcal{D}} \sum_{u'\in\mathcal{U}} \sum_{k\in\mathcal{K}} \sum_{n'=0}^{t} \sum_{a=1}^{A_{k,d'+7n'}^{i}} \mathbf{e}_{\text{u}} L_{d'+7n',k,i}^{au'}(d-d'+7(t-n'))\theta_{\mathbf{u'd'i}}^{\mathbf{k}} \\
&= \sum_{i\in\mathcal{I}} \sum_{d'\in\mathcal{D}} \sum_{u'\in\mathcal{U}} \sum_{k\in\mathcal{K}} \sum_{n=0}^{t} \sum_{a=1}^{A_{k,d'+7(t-n)}^{i}} \mathbf{e}_{\text{u}} L_{d'+7(t-n),k,i}^{au'}(d-d'+7n)\theta_{\mathbf{u'd'i}}^{\mathbf{k}}.
\end{aligned}
$$

In the first equality, we establish our stochastic census function. To do this we look at all previously admitted patients (from all units) and see if they are currently in unit u. We add up all such patients to get the census of unit u on day d of week t. The second equality above involves a simple change of variable, from n' to $n = t - n'$. This will become very useful when calculating $\mu_{\mathbf{du}}$ and $\mathbf{v_{du}}$. The derivation of the mean census is given below. For brevity we consider $r = d - d' + 7n$.

$$\mu_{\mathbf{du}} = \lim_{t\to\infty} E\left[B_{du}(t)\right] \tag{A.31}$$

$$= \lim_{t\to\infty} E\left[\sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{t}\sum_{a=1}^{A^i_{k,d'+7(t-n)}} \mathbf{e}_{\mathrm{u}} L^{au'}_{d'+7(t-n),k,i}(r)\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right] \tag{A.32}$$

$$= E\left[\sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty}\sum_{a=1}^{A^i_{k,d'+7(t-n)}} \mathbf{e}_{\mathrm{u}} L^{au'}_{d'+7(t-n),k,i}(r)\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right] \tag{A.33}$$

$$= \sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty} E\left[\sum_{a=1}^{A^i_{k,d'+7(t-n)}} \mathbf{e}_{\mathrm{u}} L^{au'}_{d'+7(t-n),k,i}(r)\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right] \tag{A.34}$$

$$= \sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty} E\left[\sum_{a=1}^{\infty} \mathbf{e}_{\mathrm{u}} L^{au'}_{d'+7(t-n),k,i}(r)\mathbb{1}\{A^i_{k,d'+7(t-n)} \geq a\}\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right] \tag{A.35}$$

$$= \sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty}\sum_{a=1}^{\infty} E\left[\mathbf{e}_{\mathrm{u}} L^{au'}_{d'+7(t-n),k,i}(r)\mathbb{1}\{A^i_{k,d'+7(t-n)} \geq a\}\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right] \tag{A.36}$$

$$= \sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty}\sum_{a=0}^{\infty} E\left[\mathbf{e}_{\mathrm{u}} L^{au'}_{d'+7(t-n),k,i}(r)\mathbb{1}\{A^i_{k,d'+7(t-n)} > a\}\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right] \tag{A.37}$$

$$= \sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty}\sum_{a=0}^{\infty} P^{u'}_{d'kui}(r)\bar{F}_{d'ki}(a)\theta^{\mathbf{k}}_{\mathbf{u'd'i}} \tag{A.38}$$

Both $\mathbf{e}_{\mathrm{u}} L^{au'}_{ski}(r)$ and $\theta^{\mathbf{k}}_{\mathbf{udi}}$ can only take values of 0 or 1, meaning $\sum_{a=1}^{A^i_{ks}} \mathbf{e}_{\mathrm{u}} L^{au'}_{ski}(r)\theta^{\mathbf{k}}_{\mathbf{udi}} \leq A^i_{ks} < \infty$. We know that $A^i_{ks} < \infty$ as there cannot be an infinite number of arrivals in one day. Thus (A.34) holds.

For (A.35), we simply replace a finite sum with an infinite sum that is equivalent via use of an indicator function, where $\mathbb{1}\{x\} = 1$ if x is true and $\mathbb{1}\{x\} = 0$ otherwise.

As $\mathbb{1}\{A^i_{ks} \geq a\}$ is binary, $\mathbf{e}_{\mathrm{u}} L^{au'}_{ski}(r)\mathbb{1}\{A^i_{ks} \geq a\}\theta^{\mathbf{k}}_{\mathbf{udi}} < \infty$, and (A.36) holds. For (A.37), we change the summation over a in order to motivate the use of the tail function of the cumulative distribution. All terms $L^{au}_{ski}(r)$ and $L^{a'u'}_{s'k'i'}(r')$ are independent of each other

provided that $(a, s, k, i) \neq (a', s', k', i')$. This is true because each $L^{au}_{ski}(r)$ term represents the care pathway of an individual patient, and each patient's care pathway is independent of all others. Also, patient care pathways and the number of arrivals in a day are independent, and these are both independent of whether we accept a given class of patient on a given day. Thus, $\mathbf{e}_{\mathrm{u}} L^{au'}_{ski}(r)$, $\mathbb{1}\{A^i_{ks} \geq a\}$, and $\theta^{\mathbf{k}}_{\mathbf{udi}}$ are mutually independent, and we can treat the expectation of their product as the product of their expectations. Also, $P^{u'}_{dkui}(r)$ and $F_{dkui}(a)$ are identical between weeks, only depending on the day of the week d. Recall that by convention, $P^{u'}_{dkui}(r) = 0$ if $r < 0$. Thus, (A.38) holds.

Proof of Theorem 2.2 The derivation of the variance of the census is given below.

$$\mathbf{v_{du}} = \lim_{t\to\infty} \operatorname{var}\left(B_{du}(t)\right) \tag{A.39}$$

$$= \lim_{t\to\infty} \operatorname{var}\left(\sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{t}\sum_{a=1}^{A^i_{k,d'+7(t-n)}} \mathbf{e}_{\mathrm{u}} L^{au'}_{d'+7(t-n),k,i}(r)\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right) \tag{A.40}$$

$$= \operatorname{var}\left(\sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty}\sum_{a=1}^{A^i_{k,d'+7(t-n)}} \mathbf{e}_{\mathrm{u}} L^{au'}_{d'+7(t-n),k,i}(r)\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right) \tag{A.41}$$

$$= \sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty} \operatorname{var}\left(\sum_{a=1}^{A^i_{k,d'+7(t-n)}} \mathbf{e}_{\mathrm{u}} L^{au'}_{d'+7(t-n),k,i}(r)\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right) \tag{A.42}$$

$$= \sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty} \operatorname{var}\left(\sum_{a=1}^{\infty} \mathbf{e}_{\mathrm{u}} L^{au'}_{d'+7(t-n),k,i}(r)\mathbb{1}\{A^i_{k,d'+7(t-n)} \geq a\}\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right) \tag{A.43}$$

$$= \sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty}\sum_{a=1}^{\infty} \operatorname{var}\left(\mathbf{e}_{\mathrm{u}} L^{au'}_{d'+7(t-n),k,i}(r)\mathbb{1}\{A^i_{k,d'+7(t-n)} \geq a\}\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right) \tag{A.44}$$

$$= \sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty}\sum_{a=1}^{\infty}\left\{E\left[\left(\mathbf{e}_{\mathrm{u}}L^{au'}_{d'+7(t-n),k,i}(r)\mathbb{1}\{A^i_{k,d'+7(t-n)}\geq a\}\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right)^2\right]\right.$$
$$\left. - E\left[\mathbf{e}_{\mathrm{u}}L^{au'}_{d'+7(t-n),k,i}(r)\mathbb{1}\{A^i_{k,d'+7(t-n)}\geq a\}\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right]^2\right\} \quad \text{(A.45)}$$

$$= \sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty}\sum_{a=0}^{\infty}\left\{E\left[\left(\mathbf{e}_{\mathrm{u}}L^{au'}_{d'+7(t-n),k,i}(r)\mathbb{1}\{A^i_{k,d'+7(t-n)}> a\}\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right)^2\right]\right.$$
$$\left. - E\left[\mathbf{e}_{\mathrm{u}}L^{au'}_{d'+7(t-n),k,i}(r)\mathbb{1}\{A^i_{k,d'+7(t-n)}> a\}\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right]^2\right\} \quad \text{(A.46)}$$

$$= \sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty}\sum_{a=0}^{\infty}\left\{P^{u'}_{d'kui}(r)\bar{F}_{d'ki}(a)\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\left[P^{u'}_{d'kui}(r)\bar{F}_{d'ki}(a)\theta^{\mathbf{k}}_{\mathbf{u'd'i}}\right]^2\right\} \quad \text{(A.47)}$$

$$= \sum_{i\in\mathcal{I}}\sum_{d'\in\mathcal{D}}\sum_{u'\in\mathcal{U}}\sum_{k\in\mathcal{K}}\sum_{n=0}^{\infty}\sum_{a=0}^{\infty}\left\{P^{u'}_{d'kui}(r)\bar{F}_{d'ki}(a)\left[P^{u'}_{d'kui}(r)\bar{F}_{d'ki}(a)\right]^2\right\}\theta^{\mathbf{k}}_{\mathbf{u'd'i}} \quad \text{(A.48)}$$

As previously discussed, each patient's care pathway is independent of other patient care pathways. Also, the number of arrivals on any given day is independent of the number of arrivals on every other day. Thus, each $\sum_{a=1}^{A^i_{ks}}\mathbf{e}_{\mathrm{u}}L^{au'}_{ski}(r)\theta^{\mathbf{k}}_{\mathbf{udi}}$ term is independent of all other similar terms, and (A.42) holds. Similar to the derivation of $\mu_{\mathbf{du}}$, we replace the finite sum with an infinite sum and an indicator function to get (A.43). Each $\mathbf{e}_{\mathrm{u}}L^{au'}_{ski}(r)\mathbb{1}\{A^i_{ks}\geq a\}\theta^{\mathbf{k}}_{\mathbf{udi}}$ term is independent of all other similar terms, and (A.44) holds. To get (A.45), we simply use the definition of variance. Similarly to the derivation of $\mu_{\mathbf{du}}$, we alter the summation over a in order to use the tail distribution, giving (A.46). As each $\mathbf{e}_{\mathrm{u}}L^{au'}_{ski}(r)\mathbb{1}\{A^i_{ks}\geq a\}\theta^{\mathbf{k}}_{\mathbf{udi}}$ term can only take a value of 0 or 1,

$$E\left[\left(\mathbf{e}_{\mathrm{u}}L^{au'}_{ski}(r)\mathbb{1}\{A^i_{ks}\geq a\}\theta^{\mathbf{k}}_{\mathbf{udi}}\right)^2\right] = E\left[\mathbf{e}_{\mathrm{u}}L^{au'}_{ski}(r)\mathbb{1}\{A^i_{ks}\geq a\}\theta^{\mathbf{k}}_{\mathbf{udi}}\right].$$

Combining this with a calculation similar to that used in our derivation of $\mu_{\mathbf{du}}$, we get (A.47). We note that $\theta^{\mathbf{k}}_{\mathbf{udi}}$ is binary, which implies that $(\theta^{\mathbf{k}}_{\mathbf{udi}})^2 = \theta^{\mathbf{k}}_{\mathbf{udi}}$. Leveraging this realization, we get (A.48).

Appendix B.

Supplements to Chapter 3

B.1. Proof of Theorem 3.2

Proof. By Conditions 1-5 in Appendix B.6, we invoke Theorem 3.1 so that

$$\sup_{z\in\mathbb{R}} |\hat{F}(z|\xi_i) - F(z|\xi_i)| \xrightarrow{p} 0 \text{ as } N \to \infty,$$

for each i. Consequently,

$$\sup_{z_1,\ldots,z_n\in\mathbb{R}} \left| \prod_{i=1}^{n} \hat{F}(z_i|\xi_i) - \prod_{i=1}^{n} F(z_i|\xi_i) \right| \xrightarrow{p} 0 \text{ as } N \to \infty \tag{B.1}$$

since

$$
\begin{aligned}
& \left|\prod_{i=1}^{n}\hat{F}(z_i|\xi_i)-\prod_{i=1}^{n}F(z_i|\xi_i)\right| \\
\leq\ & \left|\prod_{i=1}^{n}\hat{F}(z_i|\xi_i)-F(z_1|\xi_1)\prod_{i=2}^{n}\hat{F}(z_i|\xi_i)\right| + \\
& \left|F(z_1|\xi_1)\prod_{i=2}^{n}\hat{F}(z_i|\xi_i)-F(z_1|\xi_1)F(z_2|\xi_2)\prod_{i=3}^{n}\hat{F}(z_i|\xi_i)\right| + \\
& \cdots+\left|\prod_{i=1}^{n-1}F(z_i|\xi_i)\hat{F}(z_n|\xi_n)-\prod_{i=1}^{n}F(z_i|\xi_i)\right| \\
\leq\ & \sum_{i=1}^{n}|\hat{F}(z_i|\xi_i)-F(z_i|\xi_i)| \xrightarrow{p} 0
\end{aligned}
$$

by Slutsky's theorem in the last step. With the assumption that z_i is bounded and $\mathcal{X}$ is compact, we can verify that the max-plus system function f has bounded total variation uniformly over X, i.e., $\sup_{X\in\mathcal{X}}\|f(x,\cdot)\|_{TV}\leq C<\infty$ where $\|\cdot\|_{TV}$ is the total variation norm. Then

$$
\begin{aligned}
& \sup_{X\in\mathcal{X}}|\hat{E}[f(X,Z)|\xi_1,\ldots,\xi_n]-E[f(X,Z)|\xi_1,\ldots,\xi_n]| \\
& \quad=\sup_{X\in\mathcal{X}}\left|\int f(X,Z)\mathrm{d}\left(\prod_{i=1}^{n}\hat{F}(z_i|\xi_i)-\prod_{i=1}^{n}F(z_i|\xi_i)\right)\right| \\
& \quad\leq\sup_{X\in\mathcal{X}}\|f(X,\cdot)\|_{TV}\sup_{z_1,\ldots,z_n\in\mathbb{R}}\left|\prod_{i=1}^{n}\hat{F}(z_i|\xi_i)-\prod_{i=1}^{n}F(z_i|\xi_i)\right| \\
& \quad\leq C\sup_{z_1,\ldots,z_n\in\mathbb{R}}\left|\prod_{i=1}^{n}\hat{F}(z_i|\xi_i)-\prod_{i=1}^{n}F(z_i|\xi_i)\right| \\
& \quad\xrightarrow{p}0
\end{aligned}
\tag{B.2}
$$

as $N\to\infty$.

Next we argue that $\{f(X,\cdot):X\in\mathcal{X}\}$ is Glivenko-Cantelli (GC). In particular, we

express $f(X,Z)$ as $\sum_{i=1}^{n} \max\{0, \sum_{j=1}^{i} z_j - \sum_{j=1}^{i} x_j\}$. Note that the class of functions $\{\sum_{j=1}^{i} z_j - \sum_{j=1}^{i} x_j : x \in \mathcal{X}\}$ with z_j's all bounded is GC, by [137] Theorem 3. Since $\max\{0, \cdot\}$ and the addition operation are continuous, we have $f(X,Z)$ being GC. Thus, letting $\tilde{E}[\cdot|\xi_1, \ldots, \xi_n]$ be the objective of the SAA problem (3.9), we have

$$\sup_{X \in \mathcal{X}} |\tilde{E}[f(X,Z)|\xi_1, \ldots, \xi_n] - \hat{E}[f(X,Z)|\xi_1, \ldots, \xi_n]| \stackrel{a.s.}{\to} 0 \text{ as } K \to \infty. \tag{B.3}$$

Combining (C.8) and (C.9), we have

$$\sup_{X \in \mathcal{X}} |\tilde{E}[f(X,Z)|\xi_1, \ldots, \xi_n] - E[f(X,Z)|\xi_1, \ldots, \xi_n]| \stackrel{p}{\to} 0 \text{ as } N, K \to \infty.$$

Hence, using (3.10), we have

$$0 \leq H(\hat{X}^*) - H(X^*) \leq (H(\hat{X}^*) - \hat{H}(\hat{X}^*)) + (\hat{H}(\hat{X}^*) - \hat{H}(X^*)) + (\hat{H}(X^*) - H(X^*)).$$

Since $H(\hat{X}^*) - \hat{H}(\hat{X}^*)$ and $\hat{H}(X^*) - H(X^*)$ both $\stackrel{p}{\to} 0$, and $\hat{H}(\hat{X}^*) - \hat{H}(X^*) \leq 0$ by using (3.9), we have

$$0 \leq H(\hat{X}^*) - H(X^*) \leq o_p(1),$$

which implies

$$H(\hat{X}^*) - H(X^*) \stackrel{p}{\to} 0.$$

□

B.2. Proof of Theorem 3.3

Proof. For convenience, denote

$$\mathcal{U}(\gamma) = \left\{ Z : \max_{k<i\le n} \left| \frac{\sum_{j=k}^{i-1} z_j - \sum_{j=k}^{i-1} \mu_j}{\sqrt{\sum_{j=k}^{i-1} \sigma_j^2}} \right| \le \gamma \right\}$$

so that (3.18) is equivalent to $\hat{P}(Z \in \mathcal{U}(\Gamma)|\xi_1, \ldots, \xi_n) \ge 1 - \delta$. Since $\tilde{\Gamma}$ is the order statistic of K samples of $\max_{k<i\le n} \left| \frac{\sum_{j=k}^{i-1} z_j - \sum_{j=k}^{i-1} \mu_j}{\sqrt{\sum_{j=k}^{i-1} \sigma_j^2}} \right|$, and $\hat{P}(Z \in \mathcal{U}(\gamma)|\xi_1, \ldots, \xi_n)$ is the distribution function of $\max_{k<i\le n} \left| \frac{\sum_{j=k}^{i-1} z_j - \sum_{j=k}^{i-1} \mu_j}{\sqrt{\sum_{j=k}^{i-1} \sigma_j^2}} \right|$, we have $\hat{P}(Z \in \mathcal{U}(\tilde{\Gamma})|\xi_1, \ldots, \xi_n) \xrightarrow{a.s.} \hat{P}(Z \in \mathcal{U}(\Gamma)|\xi_1, \ldots, \xi_n)$ as $K \to \infty$.

Note that (3.17) is $\max_{Z\in\mathcal{U}(\Gamma)} f(X, Z)$, which dominates $\min\{q : \hat{P}(f(X, Z) \le q|\xi_1, \ldots, \xi_n) \ge 1-\delta\}$ since $\hat{P}(f(X, Z) \le \max_{Z\in\mathcal{U}(\Gamma)} f(X, Z)|\xi_1, \ldots, \xi_n) \ge \hat{P}(Z \in \mathcal{U}(\Gamma)|\xi_1, \ldots, \xi_n) \ge 1-\delta$ by our construction of $\mathcal{U}(\Gamma)$. Similarly, $\hat{H} = \max_{Z\in\mathcal{U}(\tilde{\Gamma})} f(X, Z)$ dominates $\min\{q : \hat{P}(f(X, Z) \le q|\xi_1, \ldots, \xi_n) \ge 1 - \tilde{\delta}\}$, where $1 - \tilde{\delta} = \hat{P}(Z \in \mathcal{U}(\tilde{\Gamma})|\xi_1, \ldots, \xi_n)$, which satisfies $\liminf_{K\to\infty}(1 - \tilde{\delta}) \ge 1 - \delta$ by our observation above. Hence $\liminf_{K\to\infty} \hat{H} \ge \min\{q : \hat{P}(f(X, Z) \le q|\xi_1, \ldots, \xi_n) \ge 1 - \delta\}$. Now, by an argument similar to (C.8) in the proof of Theorem 4.1, we have

$$\sup_{X\in\mathcal{X},q} |\hat{P}(f(X, Z) \le q|\xi_1, \ldots, \xi_n) - P(f(X, Z) \le q|\xi_1, \ldots, \xi_n)| \xrightarrow{p} 0 \text{ as } N \to \infty$$

This gives

$$\liminf_{K\to\infty} \hat{H} \ge \min\{q : P(f(X, Z) \le q|\xi_1, \ldots, \xi_n) \ge 1 - \delta\}$$

□

B.3. Proof of Lemma 3.1

Proof. The immediate dual formulation for $\max_{F\in\mathcal{F}} E_F\big[\max_{y\in\Omega}\sum_{i=1}^{n}(z_i - x_i)y_i\big]$ is

$$\begin{aligned} \min_{\theta,\rho} \quad & \theta + \sum_{i=1}^{n}\sum_{j=1}^{m-1} r_{ij}\rho_{ij} + \sum_{i=1}^{n}\rho_{im} \\ s.t. \quad & \theta + \sum_{i=1}^{n}\sum_{j=1}^{m} I(z_i \le q_{ij})\rho_{ij} \ge f(X,Z) \qquad Z \in \mathbb{D}, \end{aligned} \tag{B.4}$$

where θ is the dual variable corresponding to $\int_{\mathbb{D}} \mathrm{d}F(Z) = 1$. But, by analyzing the constraint, the model above can be simplified as follows

$$\theta \ge f(X,Z) - \sum_{i=1}^{n}\sum_{j=1}^{m} I(z_i \le q_{ij})\rho_{ij} \quad Z \in \mathbb{D}, \tag{B.5}$$

which is equivalent to

$$\theta \ge \max_{z\in\mathbb{D}} \max_{y\in\Omega} \left\{ \sum_{i=1}^{n}(z_i - x_i)y_i - \sum_{i=1}^{n}\sum_{j=1}^{m} I(z_i \le q_{ij})\rho_{ij} \right\} \tag{B.6}$$

$$= \max_{y\in\Omega} \max_{z\in\mathbb{D}} \left\{ \sum_{i=1}^{n}(z_i - x_i)y_i - \sum_{i=1}^{n}\sum_{j=1}^{m} I(z_i \le q_{ij})\rho_{ij} \right\} \tag{B.7}$$

$$= \max_{y\in\Omega} \max_{z\in\mathbb{D}} \left\{ \sum_{i=1}^{n} \Big((z_i - x_i)y_i - \sum_{j=1}^{m} I(z_i \le q_{ij})\rho_{ij} \Big) \right\} \tag{B.8}$$

$$= \max_{y\in\Omega} \sum_{i=1}^{n} \max_{z_i\in\mathcal{D}_i} \left\{ (z_i - x_i)y_i - \sum_{j=1}^{m} I(z_i \le q_{ij})\rho_{ij} \right\} \tag{B.9}$$

$$= \max_{y\in\Omega} \sum_{i=1}^{n} \max_{j=1,\dots,m} \max_{q_{i(j-1)}\le z_i \le q_{ij}} \left\{ (z_i - x_i)y_i - \sum_{k\ge j}^{m} \rho_{ik} \right\} \tag{B.10}$$

$$= \max_{y\in\Omega} \sum_{i=1}^{n} \max_{j=1,\dots,m} \left\{ (q_{ij} - x_i)y_i - \sum_{k\ge j}^{m} \rho_{ik} \right\}, \tag{B.11}$$

where (B.7) holds by the change of the order in maximizations. (B.10) follows by defining upper and lower quantile values for any given duration, that is true $\forall z \in Z$. Finally, for any fixed j_i, (B.11) is immediate. Therefore, we have a lower bound on θ and since there is no other limits on θ, the minimization problem (B.4) can be stated as

$$\min_{\rho} \max_{y \in \Omega} \sum_{i=1}^{n} \max_{j=1,\dots,m} \left\{ (q_{ij} - x_i) y_i - \sum_{k \geq j}^{m} \rho_{ik} \right\} + \sum_{i=1}^{n} \sum_{j=1}^{m-1} r_{ij} \rho_{ij} + \sum_{i=1}^{n} \rho_{im}.$$

□

B.4. Proof of Proposition 3.1

Proof. It can be shown that for extreme points in Ω, y_n equals to 0 or $\phi > 0$ and also y_i equals to either 0 or $y_{i+1} + 1$ for $i \leq n - 1$ [146, 147]. Hence, by noting the structure of the constraints in Ω, recursive application of the given relations for some $o = 1, ..., n + 1$ results in

$$y_i = \begin{cases} o - i, & 1 \leq i \leq o \leq n; \\ \\ n + \phi - i, & 1 \leq i \leq n, o = n + 1 \end{cases}. \tag{B.12}$$

Given the recursive structure above, we can partition the integers $1, ..., n+1$ into intervals such that $i \in [g, o]$ if and only if $i = o$ ($y_i = \pi_{io}$ for $i \in [g, o]$), which generates a one-to-one mapping of Ω's extreme points and a partition of the integers. Now, by introducing a binary variable t_{go} indicating whether $[g, o]$ is one of the partitions in $[1, n + 1]$,

$$\max_{y \in \Omega} \sum_{i=1}^{n} \max_{j=1,\dots,m} \left\{ (q_{ij} - x_i) y_i - \sum_{k \geq j}^{m} \rho_{ik} \right\}$$

can be reformulated as

$$\begin{aligned}
\max_{t} \quad & \sum_{g=1}^{n+1}\sum_{o=g}^{n+1}\left(\sum_{i=g}^{o}\max_{j=1,...,m}\left\{(q_{ij}-x_i)\pi_{io}-\sum_{k\geq j}^{m}\rho_{ik}\right\}\right)t_{go} \\
s.t. \quad & \sum_{g=1}^{n+1}\sum_{o=g}^{n+1} t_{go}=1 && \forall i\in\{1,...,n+1\} \\
& t_{go}\in\{0,1\} && 1\leq g\leq o\leq n+1,
\end{aligned} \tag{B.13}$$

where $(q_{j_{n+1}}-x_{n+1})\pi_{(n+1)o}-\sum_{k\geq j_{n+1}}^{m}\rho_{(n+1)k}=\pi_{(n+1)(n+1)}=0$. (B.13), due to the unimodularity of its constraint set, has a linear programming relaxation with an equivalent optimal objective value and binary optimal solution. Consequently, the dual of the linear relaxation of problem (B.13) can be written as

$$\begin{aligned}
\min_{\lambda} \quad & \sum_{i=1}^{n}\lambda_i \\
s.t. \quad & \sum_{i=g}^{min\{o,n\}}\lambda_i\geq\sum_{i=g}^{min\{o,n\}}\max_{j=1,...,m}\left\{(q_{ij}-x_i)\pi_{io}-\sum_{k\geq j}^{m}\rho_{ik}\right\} && 1\leq g\leq n, g\leq o\leq n+1,
\end{aligned} \tag{B.14}$$

where $(\lambda_1,...,\lambda_n)$ are the dual variables. Now, by incorporating (B.14) in problem (3.25), we have

$$\begin{aligned}
\min_{x,\rho,\lambda} \quad & \sum_{i=1}^{n}\lambda_i+\sum_{i=1}^{n}\sum_{j=1}^{m-1}r_{ij}\rho_{ij}+\sum_{i=1}^{n}\rho_{im} \\
s.t. \quad & \sum_{i=g}^{min\{o,n\}}\max_{j=1,...,m}\left\{(q_{ij}-x_i)\pi_{io}-\sum_{k\geq j}^{m}\rho_{ik}-\lambda_i\right\}\leq 0 && 1\leq g\leq n, g\leq o\leq n+1 \\
& \sum_{i=1}^{n}x_i\leq T \\
& x_i\geq 0 && \forall i\in\{1,...,n\},
\end{aligned} \tag{B.15}$$

which, by introducing $\mu_{io} = \max_{j=1,\ldots,m}\{q_{ij}\pi_{io} - \sum_{k\geq j}^{m} \rho_{ik}\}$, completes the transformation of problem (3.22) to its linear program equivalent given in (3.27). Refer to Proposition 2 in [92] for a detailed discussion. □

B.5. Proof of Theorem 3.4

Proof. Consider $\mathcal{U}$ in (3.21) and

$$\mathcal{U}_0 = \{P : P(z_i \leq q_{ij}) = s_{ij}, i = 1, \ldots, n, j = 1, \ldots, m\}$$

where s_{ij} are now calibrated so that $F(q_{ij}|\xi_i) = s_{ij}$, i.e., s_{ij} is the true conditional distribution at q_{ij} for patient i.

Since the true joint distribution $\prod_{i=1}^{n} F(z_i|\xi_i)$ lies in $\mathcal{U}_0$, we have

$$E[f(X, Z)] \leq \max_{P\in\mathcal{U}_0} E_P[f(X, Z)]$$

and hence

$$H^* = \min_{X\in\mathcal{X}} H(X) \leq \min_{X\in\mathcal{X}} \max_{P\in\mathcal{U}_0} E_P[f(X, Z)] \tag{B.16}$$

Now, under Conditions 1-5 in Appendix B.6, we invoke Theorem 3.1 to obtain $r_{ij} \xrightarrow{p} s_{ij}$ as $N \to \infty$. Next, we look at the dual of $\max_{P\in\mathcal{U}_0} E_P[f(X, Z)]$ and $\max_{P\in\mathcal{U}} E_P[f(X, Z)]$, which is

$$\min_{\rho_{ij}, j=1,\ldots,m} \sum_{i,j} s_{ij}\rho_{ij} \quad \text{subject to} \quad \sum_{i,j} I(z_i \leq q_{ij})\rho_{ij} \geq f(X, Z)\ \forall Z \tag{B.17}$$

and

$$\min_{\rho_{ij}, j=1,\ldots,m} \sum_{i,j} r_{ij}\rho_{ij} \quad \text{subject to} \quad \sum_{i,j} I(z_i \leq q_{ij})\rho_{ij} \geq f(X, Z)\ \forall Z \tag{B.18}$$

respectively, which are upper bounds to $\max_{P\in\mathcal{U}_0} E_P[f(X, Z)]$ and $\max_{P\in\mathcal{U}} E_P[f(X, Z)]$ by

weak duality. Hence

$$\begin{aligned}
&\min_{X\in\mathcal{X}} \max_{P\in\mathcal{U}_0} E_P[f(X,Z)] \\
&\leq \min_{X\in\mathcal{X},\rho_{ij},j=1,\ldots,m} \left\{ \sum_{i,j} s_{ij}\rho_{ij} : \sum_{i,j} I(z_i \leq q_{ij})\rho_{ij} \geq f(X,Z)\ \forall Z \right\} \\
&= \min_{X\in\mathcal{X},\rho_{ij},j=1,\ldots,m} \left\{ \sum_{i,j} r_{ij}\rho_{ij} : \sum_{i,j} I(z_i \leq q_{ij})\rho_{ij} \geq f(X,Z)\ \forall Z \right\} + \sum_{i,j}(s_{ij}-r_{ij})\rho_{ij}^{\tilde{\epsilon}} + \tilde{\epsilon}
\end{aligned} \tag{B.19}$$

where $(\rho_{ij}^{\tilde{\epsilon}})_{ij}$ is an $\tilde{\epsilon}$-optimal dual solution to (B.17). Note that (B.19) is equal to

$$\hat{H} + \sum_{i,j}(s_{ij}-r_{ij})\rho_{ij}^{\tilde{\epsilon}} + \tilde{\epsilon} \tag{B.20}$$

since (B.18) is equivalent to (3.27) as they both dualize $\max_{P\in\mathcal{U}} E_P[f(X,Z)]$ (but different representations), Now, we claim that we can find a $(\rho_{ij}^{\tilde{\epsilon}})_{ij}$ that is bounded. Suppose, on the contrary, that some component of $(\rho_{ij}^{\tilde{\epsilon}})_{ij}$ is infinite. Then, since s_{ij}'s and $f(X,Z)$ are all bounded, and s_{ij}'s are non-negative, we can decrease that component of $(\rho_{ij}^{\tilde{\epsilon}})_{ij}$ without increasing the objective value. Thus, combining this with (B.20), and using our deduction that $r_{ij} \xrightarrow{p} s_{ij}$ as $N \to \infty$ and (B.16), we have

$$P(H^* \leq \hat{H} + \epsilon) \to 1$$

as $N \to \infty$, for any $\epsilon > \tilde{\epsilon}$. □

B.6. QRF Consistency Conditions

Let Ξ be the predictor variable with dimensionality p and $\mathcal{B}$ be the space in which Ξ lives. Then, for the reader's convenience, the following list of assumptions is quoted from [96]:

Condition 2. *"$\mathcal{B} = [0,1]^p$ and Ξ is uniform on $[0,1]^p$."*

Alternatively, one may assume that the density of covariates is bounded from above and below by positive constants.

Condition 3. *"The proportion of observations in a node, relative to all observations, is vanishing as the size of data increases, ($n \to \infty$). The minimum number of observations in a node is non-decreasing in the limit as the size of data goes to infinity."*

Condition 4. *"The probability that a variable is chosen for the splitpoint is bounded from below for every node by a positive constant. If a node is split, the split is chosen so that each of the resulting sub-nodes contains at least a proportion γ of the observations in the original node, for some $0 < \gamma \leq 0.5$."*

Condition 5. *"There exists a constant L so that $F(z|\Xi = \xi)$ is Lipschitz continuous with parameter L, that is for all ξ, $\xi' \in \mathcal{B}$,*

$$\sup_y |F(z|\Xi = \xi) - F(z|\Xi = \xi')| \leq L||\xi - \xi'||_1."$$

Condition 6. *"The conditional distribution function $F(z|\Xi = \xi)$ is, for every $\xi \in \mathcal{B}$, strictly monotonically increasing in z."*

B.7. Formal Description of Random Forests and Quantile Regression Forests

We follow the notation of [26] for a given dataset of $((\xi_1, z_1), ..., (\xi_n, z_n))$. Consider Ξ as the predictor variable with dimensionality p and $\mathcal{B}$ as the space in which Ξ lives. Let θ be the random parameter that defines which features are selected for split-points at each

node in a tree. Then, $T(\theta)$ denotes the corresponding tree. Every leaf $l = 1, ..., L$ of a tree associates with a rectangular subspace of $\mathcal{B}$ which is denoted by R_l and every $\xi \in \mathcal{B}$ belongs to one and only one leaf which is denoted by $l(\xi, \theta)$.

For a new observation $\Xi = \xi$, $T(\theta)$ makes prediction by taking the average over the observed response values in leaf $l(\xi, \theta)$. Define the weight vector $w_i(\xi, \theta)$ such that

$$w_i(\xi, \theta) = \frac{1_{\{\xi_i \in R_{l(\xi,\theta)}\}}}{|\{\xi_j \in R_{l(\xi,\theta)}, \forall j\}|}. \tag{B.21}$$

Therefore, the prediction of a single tree is given by

$$\hat{\mu}_{tree}(\xi) = \sum_{i=1}^{n} w_i(\xi, \theta) z_i. \tag{B.22}$$

Then, for a Random forests growing an ensemble of K trees, the prediction is given by the averaged result of K single trees using i.i.d. vector θ_t for $t = 1, ..., K$, such that

$$w_i(\xi) = \frac{1}{K} \sum_{t=1}^{K} w_i(\xi, \theta_t) \tag{B.23}$$

and

$$\hat{\mu}_{RF}(\xi) = \sum_{i=1}^{n} w_i(\xi) z_i. \tag{B.24}$$

Similar to RF that approximates the conditional mean by a weighted average over the observations of the response variable, one can approximate the conditional distribution of z by the weighted average over the observations of $1_{\{z_i \leq z\}}$ such that

$$\hat{F}(z|\Xi = \xi) = \sum_{i=1}^{n} w_i(\xi) 1_{\{z_i \leq z\}}. \tag{B.25}$$

Quantile regression forests performs a similar procedure as RF except for the very last

step. QRF, at each node, keeps the value of all observations fallen under the node and estimates the conditional distribution function via B.25 using the same weight estimation method as in RF given by B.23.

Appendix C.

Supplements to Chapter 4

C.1. Proof of Theorem 4.2

Proof. Part 1.

Note that (4.6) is equivalent to

$$P_\mu(P_{\xi|\mu}(f(x,\xi) \notin \mathcal{A}) \geq \alpha) \leq \beta \tag{C.1}$$

By the Markov inequality,

$$P_\mu(P_{\xi|\mu}(f(x,\xi) \notin \mathcal{A}) \geq \alpha) \leq \frac{E_\mu[P_{\xi|\mu}(f(x,\xi) \notin \mathcal{A})]}{\alpha}$$

So

$$\frac{E_\mu[P_{\xi|\mu}(f(x,\xi) \notin \mathcal{A})]}{\alpha} \leq \beta \tag{C.2}$$

guarantees (C.1). Note that (C.2) is equivalent to

$$E_\mu[P_{\xi|\mu}(f(x,\xi) \in \mathcal{A})] \geq 1 - \alpha\beta$$

or

$$P_{\mu,\xi}(f(x,\xi) \in \mathcal{A}) \geq 1 - \alpha\beta \tag{C.3}$$

where $P_{\mu,\xi}$ denotes the joint probability with respect to the posterior distribution of μ and the stochasticity of ξ given μ. The Monte Carlo scheme in PCSG precisely generates samples according to $P_{\mu,\xi}$. Theorem 4.1 implies that choosing N satisfying (4.8) in the sampled program (4.7) to obtain x guarantees (C.3), and consequently (4.6), with probability at least $1 - \delta$. This concludes the proof of Part 1.

Part 2. We can write (4.9) as

$$P_{MC}(P_{\mu}(P_{\xi|\mu}(f(x;\xi) \in \mathcal{A}) < 1 - \alpha) \leq \beta) \geq 1 - \delta$$

This implies, with probability $p \geq 1 - \delta$, the event

$$P_{\mu}(P_{\xi|\mu}(f(x;\xi) \in \mathcal{A}) < 1 - \alpha) \leq \beta$$

occurs (under P_{MC}). And with probability $1 - p$, we have the obvious bound

$$P_{\mu}(P_{\xi|\mu}(f(x;\xi) \in \mathcal{A}) < 1 - \alpha) \leq 1$$

The last two bounds imply that

$$\begin{aligned} & E_{MC}[P_\mu(P_{\xi|\mu}(f(x;\xi) \in \mathcal{A}) < 1-\alpha)] \\ \leq\ & p\beta + (1-p) \\ \leq\ & (1-\delta)\beta + \delta \ \ \text{(by optimizing over } p \geq 1-\delta) \\ =\ & \beta + \delta - \beta\delta \\ =\ & 1 - (1-\beta)(1-\delta) \end{aligned}$$

which gives (4.10). □

C.2. Proof of Theorem 4.3

Proof. From Theorem 4.2 part 2, if we use β_t, δ_t and N_t satisfying (4.12), we have

$$E_{MC}[P_{\mu_t}(P_{\xi_t|\mu_t}(f(x;\xi_t) \in \mathcal{A}_t|\mathcal{F}_{t-1}) < 1-\alpha)] \leq 1 - (1-\beta_t)(1-\delta_t) \tag{C.4}$$

We want to show (4.14), which is equivalent to

$$P_{MC}(P_{\mu_{1:T}}(P_{\xi_t|\mu_t}(f_t(x_t,\xi_t) \in \mathcal{A}_t|\mathcal{F}_{t-1}) < 1-\alpha, \text{ for some } t \in \mathcal{S}) > \beta) \leq \lambda \tag{C.5}$$

Note that the left hand side of (C.5) is bounded from above as

$$
\begin{aligned}
& P_{MC}(P_{\mu_{1:T}}(P_{\xi_t|\mu_t}(f_t(x_t,\xi_t)\in\mathcal{A}_t|\mathcal{F}_{t-1})<1-\alpha,\ \text{for some } t\in\mathcal{S})>\beta) \\
\leq\ & P_{MC}\left(\sum_{t\in\mathcal{S}} P_{\mu_{1:T}}(P_{\xi_t|\mu_t}(f_t(x_t,\xi_t)\in\mathcal{A}_t|\mathcal{F}_{t-1})<1-\alpha)>\beta\right) \\
\leq\ & \frac{1}{\beta}E_{MC}\left[\sum_{t\in\mathcal{S}} P_{\mu_{1:T}}(P_{\xi_t|\mu_t}(f_t(x_t,\xi_t)\in\mathcal{A}_t|\mathcal{F}_{t-1})<1-\alpha)\right]\ \text{(by the Markov inequality)} \\
=\ & \frac{1}{\beta}\sum_{t\in\mathcal{S}} E_{MC}[P_{\mu_t}(P_{\xi_t|\mu_t}(f_t(x_t,\xi_t)\in\mathcal{A}_t|\mathcal{F}_{t-1})<1-\alpha)] \\
\leq\ & \frac{1}{\beta}\sum_{t\in\mathcal{S}}(1-(1-\beta_t)(1-\delta_t))\ \text{(by using (C.4))}
\end{aligned}
$$

Hence

$$\frac{1}{\beta}\sum_{t\in\mathcal{S}}(1-(1-\beta_t)(1-\delta_t))\leq\lambda \tag{C.6}$$

guarantees that (C.5) holds. Noting that (C.6) is equivalent to (4.13), we conclude our theorem. □

C.3. Proof of Proposition 4.1

Proof. Let $s=|\mathcal{S}|$. Setting $\beta_t=\delta_t=\gamma$, we have $\sum_{t\in\mathcal{S}}(\beta_t+\delta_t-\beta_t\delta_t)=s(2\gamma-\gamma^2)$. We want $s(2\gamma-\gamma^2)\leq\beta\lambda$, or equivalently

$$\gamma^2-2\gamma+\beta\lambda/s\geq 0 \tag{C.7}$$

Since the left hand side of (C.7) is a convex quadratic function, (C.7) holds if and only if

$$\gamma\geq 1+\sqrt{1-\beta\lambda/s}\quad\text{ or }\quad\gamma\leq 1-\sqrt{1-\beta\lambda/s}$$

The first condition is never satisfied since γ must be ≤ 1. The second condition is valid and gives (4.14). □

C.4. Proof of Proposition 4.2

Proof. Without loss of generality, we label t as the counter in $\mathcal{S}$ for convenience (i.e., assume that $t = \zeta(t)$ by relabeling β_t and δ_t). We want $\sum_{t=1}^{\infty}(\beta_t + \delta_t - \beta_t\delta_t) \leq \beta\lambda$ holds so that (4.14) holds regardless of T. Note that

$$\sum_{t=1}^{\infty}(\beta_t + \delta_t - \beta_t\delta_t) = 2\sum_{t=1}^{\infty}\gamma_t - \sum_{t=1}^{\infty}\gamma_t^2 \tag{C.8}$$

We analyze the two terms of the right hand side of (C.8). By definition $\gamma_t = \eta$ if and only if $t \leq 1/\eta^{1/\rho}$. Thus, for the first term, we have

$$\begin{aligned} 2\sum_{t=1}^{\infty}\gamma_t &= 2\left(\lfloor\frac{1}{\eta^{1/\rho}}\rfloor \cdot \eta + \sum_{t=\lfloor 1/\eta^{1/\rho}\rfloor+1}^{\infty}\frac{1}{t^{\rho}}\right) \\ &\leq 2\left(\eta^{1-1/\rho} + \int_{\lfloor 1/\eta^{1/\rho}\rfloor}^{\infty}\frac{1}{u^{\rho}}du\right) \\ &\leq 2\left(\eta^{1-1/\rho} + \frac{1}{\rho-1}\frac{1}{(1/\eta^{1/\rho}-1)^{\rho-1}}\right) \end{aligned} \tag{C.9}$$

For the second term in (C.8), we have

$$\begin{aligned} \sum_{t=1}^{\infty}\gamma_t^2 &= \lfloor\frac{1}{\eta^{1/\rho}}\rfloor \cdot \eta^2 + \sum_{t=\lfloor 1/\eta^{1/\rho}\rfloor+1}^{\infty}\frac{1}{t^{2\rho}} \\ &\geq \frac{\eta^2}{\eta^{1/\rho}+1} + \int_{\lfloor 1/\eta^{1/\rho}\rfloor+1}^{\infty}\frac{1}{u^{2\rho}}du \\ &\geq \frac{\eta^2}{\eta^{1/\rho}+1} + \frac{1}{2\rho-1}\frac{1}{(1/\eta^{1/\rho}+1)^{2\rho-1}} \end{aligned} \tag{C.10}$$

Therefore, combining (C.9) and (C.10) into (C.8), we have (4.15) implies $2\sum_{t=1}^{\infty}\gamma_t - \sum_{t=1}^{\infty}\gamma_t^2 \leq \beta\lambda$ or (4.14). □

C.5. Proof of Corollary 4.4

Proof. The corollary follows immediately by noticing that the linear program (4.16) is trivially convex, and applying Theorem 4.3. □

Bibliography

[1] S. Agrawal and N. Goyal. "Analysis of Thompson sampling for the multi-armed bandit problem". In: *Conference on Learning Theory*. 2012, pp. 39–1.

[2] L. Aiken, S. Clarke, D. Sloane, J. Sochalski, and J. Silber. "Hospital nurse staffing and patient mortality, nurse burnout, and job dissatisfaction". In: *JAMA* 288.16 (2002), pp. 1987–1993.

[3] K. Amin, M. Kearns, P. Key, and A. Schwaighofer. "Budget optimization for sponsored search: Censored learning in MDPs". In: *arXiv preprint arXiv:1210.4847* (2012).

[4] J. E. Anderson and D. C. Chang. "Using electronic health records for surgical quality improvement in the era of big data". In: *JAMA surgery* 150.1 (2015), pp. 24–29.

[5] E. Ang, S. Kwasnick, M. Bayati, E. L. Plambeck, and M. Aratow. "Accurate ED Wait Time Prediction". In: *Manufacturing and Service Operations Management* (Forthcoming).

[6] M. Armony, C. Chan, and B. Zhu. "Critical care in hospitals: When to introduce a Step Down Unit?" In: *Productions and Operations Management* (forthcoming 2018).

[7] J.-Y. Audibert and S. Bubeck. "Best arm identification in multi-armed bandits". In: *COLT-23th Conference on Learning Theory-2010*. 2010, 13–p.

[8] P. Auer, N. Cesa-Bianchi, and P. Fischer. "Finite-time analysis of the multiarmed bandit problem". In: *Machine learning* 47.2-3 (2002), pp. 235–256.

[9] A. Badanidiyuru, R. Kleinberg, and A. Slivkins. "Bandits with knapsacks". In: *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*. IEEE. 2013, pp. 207–216.

[10] G.-Y. Ban, N. El Karoui, and A. E. Lim. "Machine learning and portfolio optimization". In: *Management Science* (2016).

[11] G.-Y. Ban and C. Rudin. "The big data newsvendor: Practical insights from machine learning". In: (2017).

[12] C. Bandi and D. Bertsimas. "Tractable stochastic analysis in high dimensions via robust optimization". In: *Mathematical programming* (2012), pp. 1–48.

[13] J. Barado, J. Guergué, L. Esparza, C. Azcárate, F. Mallor, and S. Ochoa. "A mathematical model for simulating daily bed occupancy in an intensive care unit". In: *Critical Care Medicine* 40.4 (2012), pp. 1098–1104.

[14] H. Bastani and M. Bayati. "Online decision-making with high-dimensional covariates". In: (2015).

[15] A. Ben-Tal, D. Den Hertog, A. De Waegenaere, B. Melenberg, and G. Rennen. "Robust solutions of optimization problems affected by uncertain probabilities". In: *Management Science* 59.2 (2013), pp. 341–357.

[16] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust optimization*. Princeton University Press, 2009.

[17] A. Ben-Tal and A. Nemirovski. "Robust solutions of linear programming problems contaminated with uncertain data". In: *Mathematical programming* 88.3 (2000), pp. 411–424.

[18] D. Bertsimas, D. B. Brown, and C. Caramanis. "Theory and applications of robust optimization". In: *SIAM review* 53.3 (2011), pp. 464–501.

[19] D. Bertsimas, V. Gupta, and N. Kallus. "Data-driven robust optimization". In: *arXiv preprint arXiv:1401.0212* (2013).

[20] D. Bertsimas and N. Kallus. "From predictive to prescriptive analytics". In: *arXiv preprint arXiv:1402.5481* (2014).

[21] O. Besbes and A. Zeevi. "Blind network revenue management". In: *Operations research* 60.6 (2012), pp. 1537–1550.

[22] Y. U. Bing-Hua. "Delayed admission to intensive care unit for critically surgical patients is associated with increased mortality". In: *The American Journal of Surgery* 208.2 (2014), pp. 268–274.

[23] J. Boesel, B. L. Nelson, and S.-H. Kim. "Using ranking and selection to clean up after simulation optimization". In: *Operations Research* 51.5 (2003), pp. 814–825.

[24] L. Bouarfa, A. Schneider, H. Feussner, N. Navab, H. U. Lemke, P. P. Jonker, and J. Dankelman. "Prediction of intraoperative complexity from preoperative patient data for laparoscopic cholecystectomy". In: *Artificial intelligence in medicine* 52.3 (2011), pp. 169–176.

[25] L. Breiman, J. Friedman, C. Stone, and R. Olshen. *Classification and Regression Trees*. The Wadsworth and Brooks-Cole statistics-probability series. Taylor & Francis, 1984. ISBN: 9780412048418. URL: `https://books.google.com/books?id=JwQx-WOmSyQC`.

[26] L. Breiman. "Random forests". In: *Machine learning* 45.1 (2001), pp. 5–32.

[27] P. Bühlmann, B. Yu, et al. "Analyzing bagging". In: *The Annals of Statistics* 30.4 (2002), pp. 927–961.

[28] G. C. Calafiore. "Repetitive scenario design". In: *IEEE Transactions on Automatic Control* 62.3 (2017), pp. 1125–1137.

[29] G. C. Calafiore and M. C. Campi. "The scenario approach to robust control design". In: *IEEE Transactions on Automatic Control* 51.5 (2006), pp. 742–753.

[30] G. C. Calafiore and L. El Ghaoui. "On distributionally robust chance-constrained linear programs". In: *Journal of Optimization Theory and Applications* 130.1 (2006), pp. 1–22.

[31] G. Calafiore and M. C. Campi. "Uncertain convex programs: randomized solutions and confidence levels". In: *Mathematical Programming* 102.1 (2005), pp. 25–46.

[32] M. C. Campi and S. Garatti. "A sampling-and-discarding approach to chance-constrained optimization: feasibility and optimality". In: *Journal of Optimization Theory and Applications* 148.2 (2011), pp. 257–280.

[33] M. C. Campi and S. Garatti. "The exact feasibility of randomized solutions of uncertain convex programs". In: *SIAM Journal on Optimization* 19.3 (2008), pp. 1211–1230.

[34] A. Carè, S. Garatti, and M. C. Campi. "FAST-fast algorithm for the scenario technique". In: *Operations Research* 62.3 (2014), pp. 662–671.

[35] J. Castaing, A. Cohn, B. T. Denton, and A. Weizer. "A stochastic programming approach to reduce patient wait times and overtime in an outpatient infusion center". In: *IIE Transactions on Healthcare Systems Engineering* 6.3 (2016), pp. 111–125.

[36] P. Cerrito and J. C. Cerrito. "Data and text mining the electronic medical record to improve care and to lower costs". In: *Proceedings of SUGI*. Vol. 31. 2006, pp. 26–29.

[37] D. B. Chalfin, S. Trzeciak, A. Likourezos, B. M. Baumann, R. P. Dellinger, D.-E. study group, et al. "Impact of delayed transfer of critically ill patients from the emergency department to the intensive care unit". In: *Critical Care Medicine* 35.6 (2007), pp. 1477–1483.

[38] M. Chamanbaz, F. Dabbene, R. Tempo, V. Venkataramanan, and Q.-G. Wang. "Sequential randomized algorithms for convex optimization in the presence of uncertainty". In: *IEEE Transactions on Automatic Control* 61.9 (2016), pp. 2565–2571.

[39] C. W. Chan, L. V. Green, L. Lu, S. Lekwijit, and G. Escobar. "Assessing the Impact of Service Level when Customer Needs are Uncertain: An Empirical Investigation of Hospital Step-Down Units". In: *Management Science* (forthcoming 2018).

[40] C. W. Chan, V. F. Farias, N. Bambos, and G. J. Escobar. "Optimizing intensive care unit discharge decisions with patient readmissions". In: *Operations Research* 60.6 (2012), pp. 1323–1341.

[41] A. Charnes, W. W. Cooper, and G. H. Symonds. "Cost horizons and certainty equivalents: an approach to stochastic programming of heating oil". In: *Management Science* 4.3 (1958), pp. 235–263.

[42] L. M. Chen, M. Render, A. Sales, E. H. Kennedy, W. Wiitala, and T. P. Hofer. "Intensive care unit admitting patterns in the Veterans Affairs health care system". In: *Archives of Internal Medicine* 172.16 (2012), pp. 1220–1226.

[43] C. Combes, N. Meskens, C. Rivat, and J.-P. Vandamme. "Using a KDD process to forecast the duration of surgery". In: *International Journal of Production Economics* 112.1 (2008), pp. 279–293.

[44] M. E. Cowen, J. L. Czerwinski, P. J. Posa, E. Van Hoek, J. Mattimore, L. K. Halasyamani, and R. L. Strawderman. "Implementation of a mortality prediction rule for real-time decision making: Feasibility and validity". In: *Hospital Medicine* 9.11 (2014), pp. 720–726.

[45] M. Cowen, R. Strawderman, J. Czerwinski, M. Smith, and L. Halasyamani. "Mortality predictions on admission as a context for organizing care activities". In: *Hospital Medicine* 8.5 (2013), pp. 229–235.

[46] S. I. Davies. "Machine learning at the operating room of the future: A comparison of machine learning techniques applied to operating room scheduling". PhD thesis. Massachusetts Institute of Technology, 2004.

[47] D. P. De Farias and B. Van Roy. "On constraint sampling in the linear programming approach to approximate dynamic programming". In: *Mathematics of operations research* 29.3 (2004), pp. 462–478.

[48] B. Denton, J. Viapiano, and A. Vogl. "Optimization of surgery sequencing and scheduling decisions under uncertainty". In: *Health care management science* 10.1 (2007), pp. 13–24.

[49] B. Denton, J. Viapiano, and A. Vogl. "Optimization of surgery sequencing and scheduling decisions under uncertainty". In: *Health care management science* 10.1 (2007), pp. 13–24.

[50] F. Dexter, E. U. Dexter, D. Masursky, and N. A. Nussmeier. "Systematic review of general thoracic surgery articles to identify predictors of operating room case durations". In: *Anesthesia & Analgesia* 106.4 (2008), pp. 1232–1241.

[51] F. Dexter, J. Ledolter, V. Tiwari, and R. H. Epstein. "Value of a scheduled duration quantified in terms of equivalent numbers of historical cases". In: *Anesthesia & Analgesia* 117.1 (2013), pp. 205–210.

[52] F. Dexter and A. Macario. "Decrease in case duration required to complete an additional case during regularly scheduled hours in an operating room suite: a computer simulation study". In: *Anesthesia & Analgesia* 88.1 (1999), pp. 72–76.

[53] F. Dexter, A. Macario, and J. Ledolter. "Identification of systematic underestimation (bias) of case durations during case scheduling would not markedly reduce overutilized operating room time". In: *Journal of clinical anesthesia* 19.3 (2007), pp. 198–203.

[54] W. Ding, T. Qin, X.-D. Zhang, and T.-Y. Liu. "Multi-Armed Bandit with Budget Constraint and Variable Costs." In: AAAI. 2013.

[55] B. Efron. "Estimation and accuracy after model selection". In: *Journal of the American Statistical Association* 109.507 (2014), pp. 991–1007.

[56] M. J. Eijkemans, M. van Houdenhoven, T. Nguyen, E. Boersma, E. W. Steyerberg, and G. Kazemier. "Predicting the UnpredictableA New Prediction Model for Operating Room Times Using Individual Characteristics and the Surgeon's Estimate". In: *The Journal of the American Society of Anesthesiologists* 112.1 (2010), pp. 41–49.

[57] G. J. Escobar, J. D. Greene, M. N. Gardner, G. P. Marelich, B. Quick, and P. Kipnis. "Intra-hospital transfers to a higher level of care: Contribution to total hospital and intensive care unit (ICU) mortality and length of stay (LOS)". In: *Hospital Medicine* 6.2 (2011), pp. 74–80.

[58] K. J. Ferreira, D. Simchi-Levi, and H. Wang. "Online network revenue management using Thompson sampling". In: *Working paper* (2016).

[59] J. H. Friedman. "Greedy function approximation: a gradient boosting machine". In: *Annals of statistics* (2001), pp. 1189–1232.

[60] J. Friedman, T. Hastie, and R. Tibshirani. "Regularization paths for generalized linear models via coordinate descent". In: *Journal of statistical software* 33.1 (2010), p. 1.

[61] D. Goldfarb and G. Iyengar. "Robust portfolio selection problems". In: *Mathematics of Operations Research* 28.1 (2003), pp. 1–38.

[62] L. V. Green. "How many hospital beds?" In: *INQUIRY: The Journal of Health Care Organization, Provision, and Financing* 39.4 (2002), pp. 400–412.

[63] N. Halpern. "Can the costs of critical care be controlled?" In: *Current Opinion in Critical Care* 15 (2009), pp. 591–596.

[64] N. Halpern and S. Pastores. "Critical Care medicine in the United States 2000-2005: An analysis of bed numbers, occupancy rates, payer mix, and costs". In: *Critical Care Medicine* 38.1 (2010), pp. 65–71.

[65] E. Hans, G. Wullink, M. Van Houdenhoven, and G. Kazemier. "Robust surgery loading". In: *European Journal of Operational Research* 185.3 (2008), pp. 1038–1050.

[66] B. Hazlehurst, J. Mullooly, A. Naleway, and B. Crane. "Detecting possible vaccination reactions in clinical notes." In: *AMIA*. 2005.

[67] J. E. Helm, S. AhmadBeygi, and M. P. Van Oyen. "Design and analysis of hospital admission control for operational effectiveness". In: *Production and Operations Management* 20.3 (2011), pp. 359–374.

[68] J. E. Helm and M. P. Van Oyen. "Design and optimization methods for elective hospital admissions". In: *Operations Research* 62.6 (2014), pp. 1265–1282.

[69] L. J. Hong, Z. Huang, and H. Lam. "Learning-based Robust Optimization: Procedures and Statistical Guarantees". In: *arXiv preprint arXiv:1704.04342* (2017).

[70] W. C. Hsiao, P. Braun, D. L. Dunn, E. R. Becker, D. Yntema, D. K. Verrilli, E. Stamenovic, and S.-P. Chen. "An overview of the development and refinement of the Resource-Based Relative Value Scale: the foundation for reform of US physician payment". In: *Medical care* (1992), NS1–NS12.

[71] G. Iapichino, D. Corbella, C. Minelli, G. H. Mills, A. Artigas, D. L. Edbooke, A. Pezzi, J. Kesecioglu, N. Patroniti, M. Baras, et al. "Reasons for refusal of admission to intensive care and impact on mortality". In: *Intensive Care Medicine* 36.10 (2010), pp. 1772–1779.

[72] M. W. Isken. "Modeling and analysis of occupancy data: A healthcare capacity planning application". In: *International Journal of Information Technology & Decision Making* 1.04 (2002), pp. 707–729.

[73] A. Jebali and A. Diabat. "A stochastic model for operating room planning under capacity constraints". In: *International Journal of Production Research* 53.24 (2015), pp. 7252–7270.

[74] P. B. Jensen, L. J. Jensen, and S. Brunak. "Mining electronic health records: towards better research applications and clinical care". In: *Nature Reviews Genetics* 13.6 (2012), pp. 395–405.

[75] N. Kallus and M. Udell. "Dynamic assortment personalization in high dimensions". In: *arXiv preprint arXiv:1610.05604* (2016).

[76] E. Kayis, T. T. Khaniyev, J. Suermondt, and K. Sylvester. "A robust estimation model for surgery durations with temporal, operational, and surgery team effects". In: *Health care management science* 18.3 (2015), pp. 222–233.

[77] E. Kayis, H. Wang, M. Patel, T. Gonzalez, S. Jain, R. Ramamurthi, C. Santos, S. Singhal, J. Suermondt, and K. Sylvester. "Improving prediction of surgery duration using operational and temporal factors". In: *AMIA Annual Symposium Proceedings*. Vol. 2012. American Medical Informatics Association. 2012, p. 456.

[78] D. KC and C. Terwiesch. "An econometric analysis of patient flows in the cardiac intensive care unit". In: *Manufacturing & Service Operations Management* 14.1 (2012), pp. 50–65.

[79] S.-H. Kim, C. W. Chan, M. Olivares, and G. Escobar. "ICU admission control: An empirical study of capacity allocation and its implication for patient outcomes". In: *Management Science* 61.1 (2014), pp. 19–38.

[80] R. Koenker. *Quantile regression*. 38. Cambridge university press, 2005.

[81] C. M. Lagoa, X. Li, and M. Sznaier. "Probabilistically constrained linear programs and risk-adjusted controller design". In: *SIAM Journal on Optimization* 15.3 (2005), pp. 938–951.

[82] T. L. Lai and H. Robbins. "Asymptotically efficient adaptive allocation rules". In: *Advances in applied mathematics* 6.1 (1985), pp. 4–22.

[83] A. Larsson. "The accuracy of surgery time estimations". In: *Production Planning & Control* 24.10-11 (2013), pp. 891–902.

[84] M. A. Lejeune and A. Ruszczynski. "An efficient trajectory method for probabilistic production-inventory-distribution problems". In: *Operations Research* 55.2 (2007), pp. 378–394.

[85] Y. Li, S. Zhang, R. F. Baugh, and J. Z. Huang. "Predicting surgical case durations using ill-conditioned CPT code matrix". In: *IIE Transactions* 42.2 (2009), pp. 121–135.

[86] Y. Lin and Y. Jeon. "Random forests and adaptive nearest neighbors". In: *Journal of the American Statistical Association* 101.474 (2006), pp. 578–590.

[87] E. Litvak and M. C. Long. "Cost and quality under managed care: Irreconcilable differences". In: *Am J Manag Care* 6.3 (2000), pp. 305–12.

[88] K. Luangkesorn and Z. Eren-Dogu. "Markov Chain Monte Carlo methods for estimating surgery duration". In: *Journal of Statistical Computation and Simulation* 86.2 (2016), pp. 262–278.

[89] J. Luedtke, S. Ahmed, and G. L. Nemhauser. "An integer programming approach for linear programs with probabilistic constraints". In: *Mathematical Programming* 122.2 (2010), pp. 247–272.

[90] A. Macario. "Is it possible to predict how long a surgery will last". In: *Medscape Anesthesiology* 108.3 (2010), pp. 681–685.

[91] A. Macario and F. Dexter. "Estimating the duration of a case when the surgeon has not recently scheduled the procedure at the surgical suite". In: *Anesthesia & Analgesia* 89.5 (1999), pp. 1241–1245.

[92] H.-Y. Mak, Y. Rong, and J. Zhang. "Appointment scheduling with limited distributional information". In: *Management Science* 61.2 (2014), pp. 316–334.

[93] K. S. Mathews and E. F. Long. "A conceptual framework for improving critical care patient flow and bed use". In: *Annals of the American Thoracic Society* 12.6 (2015), pp. 886–894.

[94] M. L. McManus, M. C. Long, A. Cooper, J. Mandell, D. M. Berwick, M. Pagano, and E. Litvak. "Variability in surgical caseload and access to intensive care services". In: *The Journal of the American Society of Anesthesiologists* 98.6 (2003), pp. 1491–1496.

[95] M. McManus, M. Long, A. Cooper, J. Mandell, D. Berwick, M. Pagano, and E. Litvak. "Variability in surgical caseload and access to intensive care services". In: *Anesthesiology* 98.6 (2003), pp. 1491–1496.

[96] N. Meinshausen. "Quantile regression forests". In: *Journal of Machine Learning Research* 7.Jun (2006), pp. 983–999.

[97] A. Meisami, M. P. V. Oyen, and H. Lam. "Uncertainty quantification on simulation analysis driven by random forests". In: *2017 Winter Simulation Conference (WSC)*. Dec. 2017, pp. 3266–3274. DOI: 10.1109/WSC.2017.8248044.

[98] A. Meisami, J. Deglise-Hawkinson, M. E. Cowen, and M. P. Van Oyen. "Data-driven optimization methodology for admission control in critical care units". In: *Health Care Management Science* (2018). ISSN: 1572-9389. DOI: 10.1007/s10729-018-9439-5. URL: https://doi.org/10.1007/s10729-018-9439-5.

[99] A. Meisami, C. Stromblad, N. Kastango, M. P. Van Oyen, R. S. Wilson, and N. R. Abu-Rustum. "Towards Individualized Care Delivery Planning: Big Data Analytics and Surgical Case Duration". In: *Working paper* (2018).

[100] B. L. Miller and H. M. Wagner. "Chance constrained programming with joint constraints". In: *Operations Research* 13.6 (1965), pp. 930–945.

[101] T. B. Murdoch and A. S. Detsky. "The inevitable application of big data to health care". In: *Jama* 309.13 (2013), pp. 1351–1352.

[102] M. Olivares, C. Terwiesch, and L. Cassorla. "Structural estimation of the newsvendor model: an application to reserving operating room time". In: *Management Science* 54.1 (2008), pp. 41–55.

[103] A. Pani, S. Raghavan, and M. Sahin. "Large-Scale Advertising Portfolio Optimization in Online Marketing". In: (2017).

[104] B. Pearce, N. Hosseini, K. Taaffe, N. Huynh, and S. Harris. "Modeling interruptions and patient flow in a preoperative hospital environment". In: *Simulation Conference (WSC), Proceedings of the 2010 Winter*. IEEE. 2010, pp. 2261–2270.

[105] A. Prékopa. "Probabilistic programming". In: *Handbooks in operations research and management science* 10 (2003), pp. 267–351.

[106] J. R. Quinlan. *C4. 5: programs for machine learning*. Elsevier, 2014.

[107] J. R. Quinlan. "Induction of decision trees". In: *Machine learning* 1.1 (1986), pp. 81–106.

[108] R. Redlich, N. Opel, D. Grotegerd, K. Dohm, D. Zaremba, C. Bürger, S. Münker, L. Mühlmann, P. Wahl, W. Heindel, et al. "Prediction of Individual Response to Electroconvulsive Therapy via Machine Learning on Structural Magnetic Resonance Imaging Data". In: *JAMA psychiatry* 73.6 (2016), pp. 557–564.

[109] G. Rosenthal, C. Sirio, L. Shepardson, D. Harper, A. Rotondi, and G. Cooper. "Use of intensive care units for patients with low severity of illness". In: *Archives of Internal Medicine* 158 (1998), pp. 1144–1151.

[110] R. Rossi, B. Hnich, S. A. Tarim, and S. Prestwich. "Confidence-based reasoning in stochastic constraint programming". In: *Artificial Intelligence* 228 (2015), pp. 129–152.

[111] R. Rossi, B. Hnich, S. A. Tarim, and S. Prestwich. "Finding (α,ϑ)-solutions via Sampled SCSP". In: *Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI, Barcelona, Spain, 16-22 July, 2011*. 2011, pp. 2172–2177.

[112] M. J. Rothman, S. I. Rothman, and J. Beals. "Development and validation of a continuous measure of patient condition using the electronic medical record". In: *Biomedical Informatics* 46.5 (2013), pp. 837–848.

[113] C. Rudin and G.-Y. Vahn. "The big data newsvendor: Practical insights from machine learning". In: (2014).

[114] D. Russo and B. Van Roy. "An information-theoretic analysis of Thompson sampling". In: *The Journal of Machine Learning Research* 17.1 (2016), pp. 2442–2471.

[115] F. Sabria and C. F. Daganzo. "Approximate expressions for queueing systems with scheduled arrivals and established service order". In: *Transportation Science* 23.3 (1989), pp. 159–165.

[116] C. B. Sankey, G. McAvay, J. M. Siner, C. L. Barsky, and S. I. Chaudhry. "Deterioration to Door Time: An Exploratory Analysis of Delays in Escalation of Care for Hospitalized Patients". In: *General Internal Medicine* (2016), pp. 1–6.

[117] Z. ShahabiKargar, S. Khanna, N. Good, A. Sattar, J. Lind, and J. O Dwyer. "Predicting Procedure Duration to Improve Scheduling of Elective Surgery". In: *Pacific Rim International Conference on Artificial Intelligence*. Springer. 2014, pp. 998–1009.

[118] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on stochastic programming: modeling and theory*. SIAM, 2009.

[119] Y. Shi, J. Zhang, and K. B. Letaief. "Optimal stochastic coordinated beamforming for wireless cooperative networks with CSI uncertainty". In: *IEEE Transactions on Signal Processing* 63.4 (2015), pp. 960–973.

[120] A. Shmueli, C. L. Sprung, and E. H. Kaplan. "Optimizing admissions to an intensive care unit". In: *Health Care Management Science* 6 (2003), pp. 131–136.

[121] A. Shmueli and C. L. Sprung. "Assessing the in-hospital survival benefits of intensive care". In: *International journal of technology assessment in health care* 21.1 (2005), pp. 66–72.

[122] E. Simchen, C. Sprung, N. Galai, Y. Zitser-Gurevich, Y. Bar-Lavi, L. Levi, F. Zveibil, M. Mandel, G. Mnatzaganian, N. Goldschmidt, A. Ekka-Zohar, and I. Weiss-Salz. "Survival of critically ill patients hospitalized in and out of intensive care". In: *Critical Care Medicine* 35.2 (2007), pp. 449–457.

[123] T. Sinuff, K. Kahnamoui, D. Cook, J. Luce, and M. Levy. "Rationing critical care beds: A systematic review". In: *Critical Care Medicine* 32.7 (2004), pp. 1588–1597.

[124] D. Sitompul and S. Randhawa. "Nurse scheduling models: a state-of-the-art review." In: *Journal of the Society for Health Systems* 2.1 (1990), pp. 62–72.

[125] M. Smith, G. Halvorson, and G. Kaplan. "What's needed is a health care system that learns: recommendations from an IOM report". In: *Jama* 308.16 (2012), pp. 1637–1638.

[126] H. Stelfox, B. Hemmelgarn, S. Bagshaw, S. Gao, C. Doig, C. Nijssen-Jordan, and B. Manns. "Intensive care unit bed availability and outcomes for hospitalized patients with sudden clinical deterioration". In: *Archives of Internal Medicine* 172.6 (2012), pp. 467–474.

[127] P. S. Stepaniak, C. Heij, and G. De Vries. "Modeling and prediction of surgical procedure times". In: *Statistica Neerlandica* 64.1 (2010), pp. 1–18.

[128] D. P. Strum, J. H. May, and L. G. Vargas. "Modeling the Uncertainty of Surgical Procedure TimesComparison of Log-normal and Normal Models". In: *The Journal of the American Society of Anesthesiologists* 92.4 (2000), pp. 1160–1167.

[129] D. P. Strum, A. R. Sampson, J. H. May, and L. G. Vargas. "Surgeon and type of anesthesia predict variability in surgical procedure times". In: *The Journal of the American Society of Anesthesiologists* 92.5 (2000), pp. 1454–1466.

[130] Task Force of the American College of Critical Care Medicine, Society of Critical Care Medicine. "Guidelines for intensive care unit admission, discharge, and triage". In: *Critical Care Medicine* 27 (1999), pp. 633–638.

[131] R. Tibshirani. "Regression shrinkage and selection via the lasso". In: *Journal of the Royal Statistical Society. Series B (Methodological)* (1996), pp. 267–288.

[132] L. Tran–Thanh, A. Chapman, A. Rogers, and N. R. Jennings. "Knapsack Based Optimal Policies for Budget-Limited Multi-Armed Bandits." In: AAAI. 2012.

[133] L. Tran-Thanh, L. Stavrogiannis, V. Naroditskiy, V. Robu, N. R. Jennings, and P. Key. "Efficient regret bounds for online bid optimisation in budget-limited sponsored search auctions". In: UAI. 2014.

[134] T. Tulabandhula and C. Rudin. "Robust optimization using machine learning for uncertainty sets". In: *arXiv preprint arXiv:1407.1097* (2014).

[135] N. T. Tung, J. Z. Huang, I. Khan, M. J. Li, and G. Williams. "Extensions to quantile regression forests for very high-dimensional data". In: *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer. 2014, pp. 247–258.

[136] US Census Bureau. Dec. 2013.

[137] A. Van Der Vaart and J. A. Wellner. "Preservation theorems for Glivenko-Cantelli and uniform Glivenko-Cantelli classes". In: *High dimensional probability II*. Springer, 2000, pp. 115–133.

[138] S. S. Villar, J. Bowden, and J. Wason. "Multi-armed bandit models for the optimal design of clinical trials: benefits and challenges". In: *Statistical science: a review journal of the Institute of Mathematical Statistics* 30.2 (2015), p. 199.

[139] K. Walter, M. Siegler, and J. Hall. "How decisions are made to admit patients to medical intensive care units (MICUs): A survey of MICU directors at academic medical centers across the United States". In: *Critical Care Medicine* 36.2 (2008), pp. 414–420.

[140] P. P. Wang. "Static and dynamic scheduling of customer arrivals to a single-server system". In: *Naval Research Logistics (NRL)* 40.3 (1993), pp. 345–360.

[141] Z. Wang, S. Deng, and Y. Ye. "Close the gaps: A learning-while-doing algorithm for single-product revenue management problems". In: *Operations Research* 62.2 (2014), pp. 318–331.

[142] W. Wiesemann, D. Kuhn, and M. Sim. "Distributionally robust convex optimization". In: *Operations Research* 62.6 (2014), pp. 1358–1376.

[143] I. H. Wright, C. Kooperberg, B. A. Bonar, and G. Bashein. "Statistical Modeling to Predict Elective Surgery TimeComparison with a Computer Scheduling System and Surgeon-provided Estimates". In: *The Journal of the American Society of Anesthesiologists* 85.6 (1996), pp. 1235–1245.

[144] Y. Xia, H. Li, T. Qin, N. Yu, and T.-Y. Liu. "Thompson Sampling for Budgeted Multi-Armed Bandits." In: IJCAI. 2015.

[145] M. Yang, M. J. Fry, and C. Scurlock. "The ICU will see you now: efficient equitable admission control policies for a surgical ICU with batch arrivals". In: *IIE Transactions* 47.6 (2015), pp. 586–599.

[146] W. I. Zangwill. "A backlogging model and a multi-echelon model of a dynamic economic lot size production system-a network approach". In: *Management Science* 15.9 (1969), pp. 506–527.

[147] W. I. Zangwill. "A deterministic multi-period production scheduling model with backlogging". In: *Management Science* 13.1 (1966), pp. 105–119.

[148] J. Zhou and F. Dexter. "Method to assist in the scheduling of add-on surgical cases-upper prediction bounds for surgical case durations based on the log-normal distribution". In: *The Journal of the American Society of Anesthesiologists* 89.5 (1998), pp. 1228–1232.

[149] J. E. Zimmerman and A. A. Kramer. "A model for identifying patients who may not need intensive care unit admission". In: *Critical Care* 25.2 (2010), pp. 205–213.

[150] S. Zymler, D. Kuhn, and B. Rustem. "Distributionally robust joint chance constraints with second-order moment information". In: *Mathematical Programming* (2013), pp. 1–32.