Title
Data-Driven Decision-Making Under Uncertainty with Applications in Healthcare and Energy Management

Author
Ghodsi, Saeed

Publication Date
2022

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Data-Driven Decision-Making Under Uncertainty with Applications in Healthcare and Energy Management

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Management

by

Saeed Ghodsi

2022

ABSTRACT OF THE DISSERTATION

Data-Driven Decision-Making
Under Uncertainty with Applications in
Healthcare and Energy Management

by

Saeed Ghodsi
Doctor of Philosophy in Management
University of California, Los Angeles, 2022
Professor Reza H Ahmadi, Chair

Decision-making under uncertainty has been studied for a long time by the operations management research community. In the past, uncertainty models were often derived based on domain knowledge. However, the availability of vast amounts of data in the recent years has shifted interests towards data-driven approaches for uncertainty quantification. More specifically, statistical models are employed within this framework for characterizing the uncertain components of a stochastic optimization problem based on historical data.

In this dissertation, we focus on applications of data-driven decision-making under uncertainty in the healthcare and energy management sectors. The first part of our work provides a mathematical framework for efficient call assignment under Direct Load Control (DLC) contracts (i.e. an incentive-based demand-response program that is widely used by utility firms for balancing the supply and demand of electricity during peak times). Specifically, we employ a model for forecasting energy consumption and develop a large-scale integer stochastic dynamic optimization problem. We then propose a novel hierarchical

approximation scheme for efficient execution of the contracts. We evaluate the quality of our proposed approach using real-world data obtained from California Independent System Operator (CAISO), which is the umbrella organization of utility firms in California. A large utility firm in California has implemented our model and informed us that they have experienced a 4% additional reduction in their cost.

Following a similar predict-then-optimize methodological framework, the second part of this dissertation studies data-driven healthcare intervention planning. Specifically, we develop a continuous-time latent-space Markovian model for describing disease progression based on discrete-time irregularly-spaced observations. Our model is capable of incorporating the effect of interventions on progression of disease. We discuss the computational challenges of parameter estimation for this model and present a novel efficient estimation approach based on the Expectation-Maximization (EM) algorithm. A population-level optimization model for intervention planning in the behavioral healthcare sector is then developed using the fitted disease progression model. Afterward, we present an extension of the model, which is more appropriate for medical healthcare domains such as cancer maintenance therapy, and formulate an EM algorithm for estimating the model parameters. Finally, we develop an individual-level intervention planning problem based on the patient's historical data using the estimated model.

The dissertation of Saeed Ghodsi is approved.

Yingnian Wu

Kumar Rajaram

Sriram Dasu

Reza H Ahmadi, Committee Chair

University of California, Los Angeles

2022

In dedication to my parents

for their endless love, support and encouragement.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

First and foremost, I would like to express my deepest appreciation to my brilliant Management Ph.D. advisor and committee chair, Professor Reza Ahmadi, from whom I learned academic integrity, critical thinking, and persistence in research work. Reza is indeed the most influential individual in my academic life and his thoughtful consultations have always helped me to find my way during turbulent times. Moreover, I am very grateful to Professor Sriram Dasu, who has played a key role in my entire academic journey, and made this dissertation possible with his excellent mentorship and guidance. He taught me how to think outside the box when confronting challenging problems and I always admire his creativity and curiosity in finding and solving fundamental research problems. Also, I would like to express my gratitude to my committee member and one of the best teachers that I've ever had, Professor Kumar Rajaram, for his insightful comments that strengthened this dissertation. Finally, I would like to offer my sincere appreciation to my Management PhD committee member and Statistics Master's advisor, Professor Yingnian Wu, for his continued support and encouragement as well as his compassionate and intuitive teaching style, which made every single class that I took with him a unique experience.

VITA

EDUCATION

University of California, Los Angeles 2018 - 2022

CPhil in *Management*; MSc in *Statistics*

Sharif University of Technology 2011 - 2018

MSc in *Electrical Engineering*; BSc in *Electrical Engineering*; Minor in *Computer Engineering*

RESEARCH EXPERIENCE

Research Scholar at University of California, Los Angeles 2018 - 2022

1- Hierarchical large-scale stochastic optimization modeling for DLCC call assignment

2- Efficient learning of CT-HMM parameters for healthcare intervention planning

3- Probabilistic social network modeling of micro-lending joint liability contracts

Research Scholar at Sharif University of Technology 2016 - 2018

1- Dynamic alignment of multi-dimensional time-series data for human activity recognition

WORK EXPERIENCE

Data and Applied Sciences Intern at Microsoft Corporation June 2022 - Sep 2022

A member of the *Connected Services and Experiences (CSX) Data R&D* team.

Data and Applied Sciences Intern at Uber Technologies, Inc. June 2021 - Sep 2021

A member of the *Core Analytics and Science* team.

PUBLICATIONS

S. Ghodsi, H. Mohammadzade, E. Korki. "Simultaneous Joint and Object Trajectory Templates for Human Activity Recognition from 3D Data." *Journal of Visual Communication and Image Representation 55* (2018)

CHAPTER 1

Introduction

In this dissertation, we study two practical problems in the domains of healthcare and energy management within the data-driven optimization framework. Chapter 2 presents a large-scale stochastic dynamic programming problem for executing Direct Load Control (DLC) contracts, and shows that the problem is NP-hard. We estimate the model parameters by constructing a demand model, and then develop a hierarchical approximation scheme for solving the optimization problem. In Chapter 3, we deign a population-level healthcare intervention planning framework for behavioral healthcare settings. As part of the planning problem, we need to estimate the transition rates between different health condition levels. To do so, we develop a disease progression model that incorporates the effect of interventions into the model, and present an efficient algorithm for estimating the parameters. We then present an extension of the model, which is more appropriate for medical healthcare domains such as cancer maintenance therapy, and study the individual-level intervention planning problem based on this model. The common theme between these two research projects is the predict-then-optimize nature of the problems. The core contribution of the first project is in efficiently solving the large-scale optimization problem, while the main challenge in the second project is in designing an efficient algorithm for fitting the prediction model. In the following, we'll present a brief overview of the next chapters.

Direct load control contracts are among the most common types of Demand Response (DR) programs that are proposed by utility firms for controlling peak energy consumption. Under DCL contracts, the customer agrees to reduce their load at certain times for a specific

period of time in response to the request of the utility firm. The contract specifies limits on the total number of times that such requests can be made as well as the total duration of load reduction on both monthly and annual basis. This research project studies the optimal assignment problem for the so-called load reduction "calls". We assume the customers are divided into a set of groups, and the utility firm wants to assign a certain number of calls (with specific beginning time and duration) to each group in a way that the total cost of energy is minimized under the contractual constraints.

We represent the problem as a large-scale stochastic dynamic programming problem in which the demand is forecasted using an auto-regressive model. We first prove that even the deterministic version of the problem is NP-hard. We then introduce a hierarchical approximation scheme for solving the problem, which includes a monthly resource allocation problem as well as a within-a-month call assignment problem. Specifically, the aggregate problem allocates a certain number of calls and hours to each month, and also determines the corresponding "threshold" value. To do so, we generate a set of samples and estimate the parameters of the optimization problem via simulation. For industrial-sized instances, even solving the aggregate problem is challenging from the computational perspective. Hence, we relax some constraints by assuming that the remaining calls and hours at the end of each month cannot be used in the consecutive months. As a result, the approximate aggregate problem can be represented in terms of the multiple-choice knapsack problem, which can be solved in a reasonable time. The monthly problem receives the total number of allocated calls and hours as well as the threshold as its input and assigns the allocated calls to the groups. The structure of this problem is similar to the original problem, and it can be shown that it is NP-hard as well. We present a heuristic approach, called scenario-based disaggregation, for approximately solving the monthly problem.

We construct a lower bound on the original problem for evaluating the performance of the proposed approximation scheme. Our numerical experiments demonstrate that the error introduced by this approximation is relatively small (i.e. on average around five percent).

Moreover, the total error is around twelve percent. We also show that the algorithm allocates calls and hours to different groups in a balanced manner. We then perform an in-depth analysis of risk and discuss how managers can employ threshold as a tool for controlling the trade-off between the average cost of energy and the level of risk. Finally, we provide an extension of the model in which the customer compliance is stochastic.

Our research on healthcare intervention planning was originally motivated by observing the trade-off between allocating the limited healthcare resources for providing care to patients as needed and an alternative approach which is providing early intervention services. Such trade-off exits in a variety healthcare settings, especially within the behavioral health sector. Therefore, we design a framework for population-level intervention planning under resource and budget constraints. For estimating the transition rates of the planning problem, we develop a statistical disease progression model that assumes the underlying health state of the patient is unobserved and employs historical data for estimating the model parameters. We discuss that parameter estimation based on traditional approaches such as Monte Carlo Expectation Maximization (MCEM) are computationally challenging for our problem, and present an efficient estimation algorithm.

Afterward, we present an extended version of the disease progression model in which the underlying health state is occasionally observed. This model is more appropriate for medical healthcare domains such as cancer maintenance therapy and heart disease management. In these contexts, an accurate examination often imposes risk or financial cost to the patient, and they are done occasionally as a result. We develop an EM algorithm for estimating the parameters of the extended model. Finally, we discuss how the extended model can be used for individual-level intervention planning based on the patient's historical data. To evaluate the performance of the estimation algorithms, we generate synthetic data and quantitatively analyze the convergence of the parameters to their true values.

CHAPTER 2

An Optimization Framework for Call Assignment Under DLC Contracts

2.1 Introduction

The demand for electricity varies over the course of a year and during a day. Power generation and distribution systems are capital intensive, and altering the installed capacity of these systems takes years. At present, available electricity storage technologies are not commercially viable at scale [Qdr06, Tem18]. Due to these factors, supply and demand for electricity have to be matched in real time.

Power-generation technologies vary considerably in generation cost, scale, and ease of altering the output [Pos15]. Peak periods are defined as time intervals in which the capacity of primary generators, such as renewable, nuclear, and gas, are not sufficient to meet demand, and utilities are forced to turn on secondary sources such as diesel or gasoline-powered generators. The problem with the deployment of secondary electricity sources is that supply costs rise sharply [Meu14, Pos15] as illustrated in Figure 2.1. According to the U.S. Government Accountability Office, generating electricity during hot summer days is 10 times more costly during summer nights. Excessive demand is a common cause of grid failure that can cost utilities billions of dollars each year.[1]

Utilities have introduced programs to influence demand patterns and reduce demand

[1] https://www.sandc.com/en/solutions/reliability

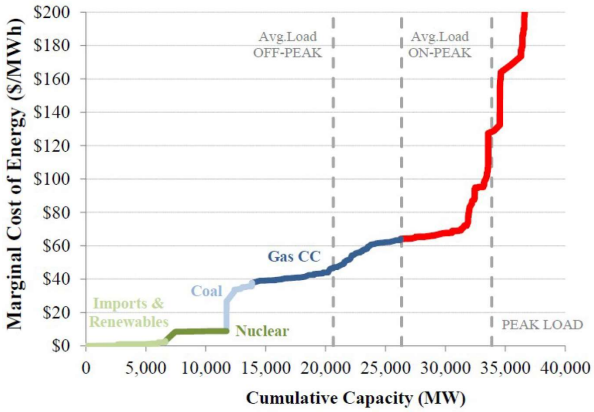

Figure 2.1: Marginal cost of energy [sei16]

during peak periods [ALW12, FKG12]. Demand response programs are an important subset of mechanisms that utilities employ to manage electricity demand [PB15, BKT06, FHT09, DG12, CKS11, AMY10, Wis11, KOK13]. For example, the Los Angeles Department of Water & Power (LADWP) offers financial incentives to businesses that voluntarily reduce their energy consumption during DR events.[2] Pacific Gas and Electric (PG&E),[3] Southern California Edison (SCE),[4] Enel X North America,[5] and many other utility firms offer similar incentive-based DR programs to their customers.

Direct load control contracts (DLCCs) are an important class of incentive-based DR programs, which permit utility companies to reduce a customer's energy usage by notifying them to reduce their consumption by a given amount based on their contracts. According to the Federal Energy Regulatory Commission, as of 2012, more than 200 utilities across the U.S. offered some type of direct load control program for residential customers [Com16]. Many utility companies, including LADWP, PG&E, and SCE, offer DLCCs. For example,

[2] https://www.ladwp.com/ladwp/faces/ladwp/commercial/c-savemoney/c-sm-rebatesandprograms/c-sm-rp-demandresponse

[3] https://www.pge.com/en_US/residential/save-energy-money/savings-solutions-and-rebates/demand-response/demand-response.page

[4] https://www.sce.com/residential/demand-response

[5] https://www.enelx.com/n-a/en/businesses/distributed-energy/demand-response

customers who participate in SCE's Time-of-Use Base Interruptible Program commit to reduce their energy usage when they receive a notification. SCE can call customers 24 hours a day, 7 days a week, and 365 days a year. However, the calls are limited to one call per day up to six hours, 10 calls per calendar month, and 180 hours per calendar year. The implementation of DLCCs is expected to rapidly expand due to the recent developments in the Internet of Things technology.

Engineering factors place constraints on how load is shed across the user base [TB11, BTM12]. As a result, utilities partition customers into a set of identical groups. The composition of each group ensures the grid engineering constraints are satisfied when the members of the group shed load. Grouping also enables the utilities to shed load in discrete increments.

In this research project, we study a class of DLCCs that are executed as follows. At the beginning of each day, the utility firm must decide, based on forecasts, which groups of customers, if any, to "call," and the starting time and the duration for each group called. A load-reduction episode is referred to as a "call." The following considerations apply while assigning the calls: (i) Customers must be informed of the call at the beginning of the day, if not earlier (some utilities send notification one day/working day in advance); (ii) the length of a call cannot exceed a predetermined value (L); (iii) a group can be called only once in a day, no more than $\breve{k}_m$ times in month m, and the total number of hours called in the month cannot exceed $\breve{h}_m$; and (iv) the annual calls are contractually limited to K calls and H hours for each customer.

Hourly demand in a day is referred to as an energy consumption profile (ECP). Figure 2.2 shows ECPs for the first Mondays of July, August, and November from 2014 to 2016. The horizontal axis denotes the time of day, and the vertical axis shows the hourly energy consumption in gigawatt hours (GWh). Note ECPs can have many different shapes; the shape of the ECPs vary by season, and, in general, each day is likely to have a unique ECP. Figure 2.3 plots the daily peak consumption for July, August, and November from 2014 to

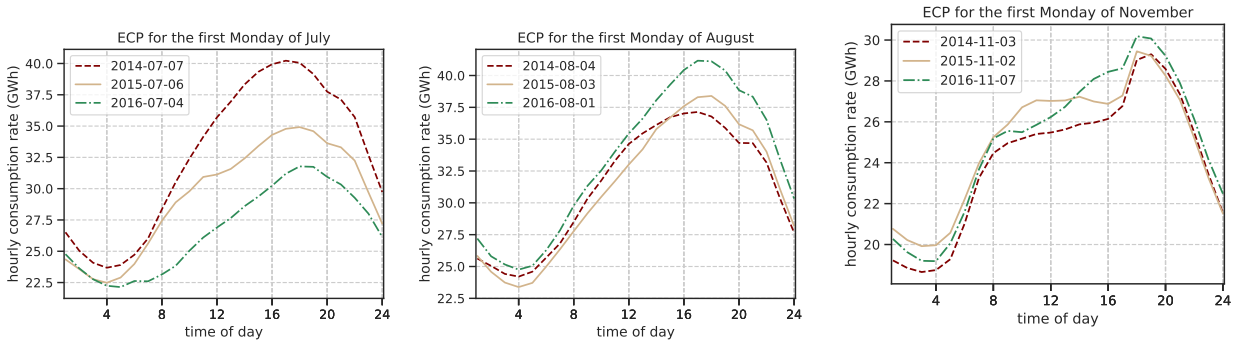

Figure 2.2: Comparison of the ECPs of the first Mondays of July, August, and November of different years (CAISO)

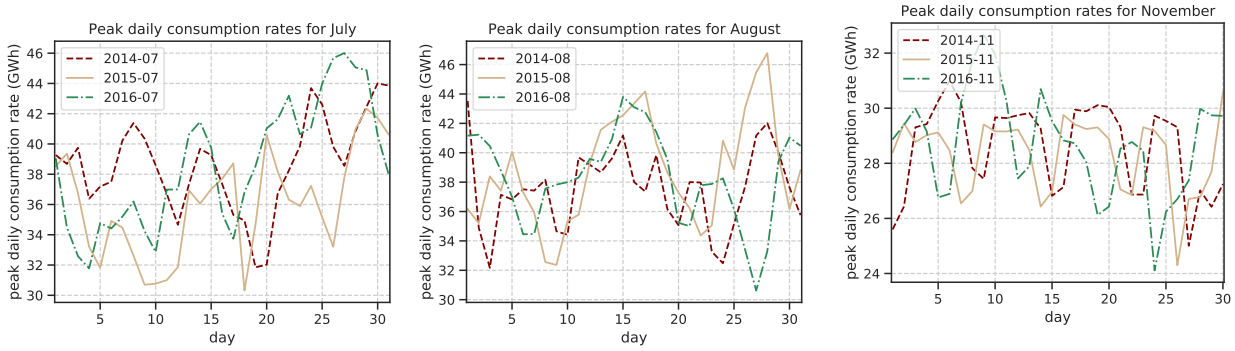

Figure 2.3: The fluctuation patterns of the peak daily consumption rate in July, August, and November of different years (CAISO)

2016. We observe variations in the peak loads within a month and across the years for the same month. Many factors such as weather patterns, hour, and weekday can affect the ECPs.

The problem of determining calls can be formulated as an integer stochastic dynamic program, where the state space includes the number of calls and hours available for each group in the current month and the rest of the year and the information needed to forecast demand over the rest of the horizon. The decision variables are calls, and the objective is to minimize the total power-generation cost over a one-year planning horizon. This stochastic dynamic program is extremely difficult, and we prove it is strongly NP-hard. Thus, we develop an approximation approach and test its performance on real data from CAISO.

We employ a hierarchical approximation approach, which consists of an annual problem and monthly problems. We incorporate a *reduce-to-threshold* policy (or *threshold* policy for short), an important managerial consideration of a major utility firm that we collaborated with. This firm requires a threshold level for each month. If the consumption is expected to exceed this threshold, it should be reduced to the threshold. This policy is motivated by the utility company's desire to diminish the need to frequently change the supply capacity. In addition, we show this policy has desirable optimality and computational properties.

We performed extensive numerical experiments using real data from CAISO to demonstrate the value of our solution approach. We found the error of our approach, measured against a lower bound, is, on average, 12.03%. More importantly, we found our solution approach can reduce cost by 3.11%, on average, and as much as 4.70%. Given the capital-intensive nature of this industry, such savings would enable utility firms to free up financial resources for future investments and offer incentives to motivate new customers to enroll in such DR programs. Additionally, we performed various sensitivity analyses related to customer preferences for the frequency and duration of calls, and obtained insights on the trade-off between reducing cost and customer preferences. Our partner utility firm notified us our model resulted in an additional 4% reduction in cost, relative to the firm's prior practice.

The remainder of the chapter is organized as follows. §2.2 provides a literature review. §2.3 formally models our integer stochastic dynamic problem, and presents our hierarchical modeling approach, which consists of a master problem, referred to as the *aggregate monthly problem* (AMP), and a subproblem, referred to as the *within-a-month problem* (WP). We also establish the computational complexity of our problem and introduce a lower bound to measure the quality of our solution approach. §2.4 and §2.5, respectively, present our analysis of AMP and WP. In §2.6, we provide the setup for our numerical experiments, and demonstrate the quality of our approximation on simulated and real data. We also describe our data collection and forecasting ECPs, and compare our approach with the prior practices. §2.7 extends our models and incorporates stochastic customer compliance behavior. §2.8

concludes the chapter.

2.2 Literature Review

Our work is related to the literature on operational issues on the energy industry and, in particular, in electric systems. Several papers in this literature focus on operational issues in the supply side of the electricity market. For example, [ACW17] study the impact of intermittent renewable power generation on supply-function competition among conventional power generators. [SB19] analyze a supply-function competition between conventional and renewable firms in a day-ahead electricity market, and prove that imposing or increasing a market-based undersupply penalty rate can lead to lower equilibrium reliability in all periods. [SS21] study the effects of rooftop solar panels on utility profits and social welfare under net metering and find that, contrary to the common belief, the net-metered distributed generation can strictly improve the utility's expected profit.

A vast and growing literature addresses various aspects of demand-side management by utility firms [VZV14]. Within demand-side management, DR instruments focus on reducing electricity consumption during peak periods, either by dynamically varying the price of electricity based on time of use or by offering incentives to customers, who reduce their consumptions when the system operator requests. For a literature review on DR programs, refer to [AA13, Sia14, DYC15, HSE16, SSS16]. Related to the dynamic pricing of electricity, [KO02] present a new contract form for the supply and procurement of interruptible electricity that bundles simple forwards with exotic call options. [BKT06] consider interruptible pay-as-you-go and pay-in-advance contracts that allow an electricity retailer to interrupt an "industrial" customer's load. They focus on deriving the optimal aggregate load interruption under each of the contracts and assume the number of industrial customers is so large that the retailer's interruption actions will not be restricted by individual constraints on the number of interruptions. [PB15] study the repeated interaction between peak producers with uncertainty

in their short-term revenues and market power and characterize the conditions under which peak producers can implicitly coordinate to achieve high prices. In their framework, the firm's demand is determined by the customer's adaptability cost and a measure of how far the ambient environment deviates from the customer's desired level.

An extensive literature on mechanism design and efficient rationing focuses on the design of incentive mechanisms that induce customers to reveal their preferences for load control options, which guide the design and practice of load-control contract. Under such mechanisms, customers can, for instance, prioritize segments of their loads or controllable devices at their premises and be compensated accordingly. [CW87], [CO16], and [PBF13] are some examples of such approaches. In this research project, we primarily focus on optimally implementing these contracts.

Utilities use DLCCs to directly control customers' electricity consumption. The literature on DLCCs is scarce [PD11, RV08, RCO09, WWG10]. [FDA22] develop an aggregation-disaggregation approximation for implementing a class of DLCCs and show it is asymptotically optimal as the length of the contract horizon grows. Our work is different because (i) we have monthly constraints on the number of calls and hours assigned to each group, and (ii) our partner utility requires that load reduction occurs such that the consumption profile is reduced to a threshold. The approach proposed by [FDA22] does not apply to DLCCs with these two features. In addition, we test the quality of our solution approach on a data from CAISO and a large utility firm in California. We show that, relative to the firm's prior practice for managing the peak load, our approximation approach can significantly reduce the firm's total cost.

Our problem that assigns calls to customer groups to flatten the ECPs has some similarities with the *online stochastic bin packing* problem with *cardinality constraints* [BBD20], in which items (calls) of unknown sizes randomly arrive and must be assigned to various bins. The number of items and the total size of the items assigned to a bin cannot exceed the capacity and the maximum number of items allowed in each bin. Our problem is different because

(i) we must determine the length of the calls (size of the items), (ii) at most one call can be assigned to a group in a day, and (iii) we have monthly limitations on the calls and hours assigned to each group.

Given the complex nature of our problem, we develop a hierarchical approximation approach that assigns calls to participating customer groups on a daily basis. We also develop a lower bounding procedure for the total cost. We test the performance of our approach using real data and show that the gap between our solution and the lower bound is reasonably small. Moreover, we report the results of our collaboration with one of the utilities in California that is currently in the process of fully implementing our approach. We show that our solution approach significantly reduces total system cost compared to the policy that is currently used by our partner utility.

2.3 Model Formulation

The problem of determining calls each day, denoted by P, can be formulated as

$$
\begin{aligned}
\text{(P)} \qquad v_{m,d}(\mathbf{X},\mathbf{Y},\mathbf{k}^{\circ}_m,\mathbf{h}^{\circ}_m,\Omega_{m,d}) \;=\; & \min\{\phi_{m,d}(\mathbf{k},\mathbf{h},\Omega_{m,d}) \\
& + \; \mathbf{E}[v_{m,d+1}(\mathbf{X}-\mathbf{k},\mathbf{Y}-\mathbf{h},\mathbf{k}^{\circ}_m-\mathbf{k},\mathbf{h}^{\circ}_m-\mathbf{h},\tilde{\Omega}_{m,d+1})]\}, && \text{(2.1a)} \\
\text{s.t.} \quad \mathbf{k} \;\leq\; & \min\{\mathbf{X},\mathbf{k}^{\circ}_m\}, && \text{(2.1b)} \\
\mathbf{h} \;\leq\; & \min\{\mathbf{Y},\mathbf{h}^{\circ}_m\}, && \text{(2.1c)}
\end{aligned}
$$

where $v_{m,d}(\cdot)$ denotes the expected future cost starting from day $d \in \{1,\ldots,\mathcal{D}_m\}$ of month $m \in \{1,\ldots,\mathcal{M}\}$ until the end of the horizon. The optimal value of P is $v_{1,1}(\mathbf{K},\mathbf{H},\breve{\mathbf{k}}_1,\breve{\mathbf{h}}_1,\Omega_{1,1})$, where vectors $\mathbf{K}$ and $\mathbf{H}$ are of size G (number of groups), with all elements equal to K and H, respectively, representing the annual limits on the calls and hours for customer groups. Similarly, all elements of vectors $\breve{\mathbf{k}}_m$ and $\breve{\mathbf{h}}_m$ are respectively equal to $\breve{k}_m$ and $\breve{h}_m$, representing the monthly contractual limits on calls and hours. Let $\Omega_{m,d}$ represent the latest information

available to forecast ECPs over the rest of the horizon. Vectors $\mathbf{X}$ and $\mathbf{Y}$ denote the number of calls and hours left for each group based on the annual contractual limits, vectors $\mathbf{k}^{\circ}_{m}$ and $\mathbf{k}^{\circ}_{m}$ denote the number of calls and hours left for each group in month m at the beginning of day d, and vectors $\mathbf{k}$ and $\mathbf{h}$ are decision variables and denote the number of calls and hours deployed on day d of month m for each group (elements of $\mathbf{k}$ are either 0 or 1). Stage cost $\phi_{m,d}(\cdot)$ is the total cost on day d of month m. The expectation in (2.1a) is with respect to future demand. Constraints (2.1b) and (2.1c) ensure the monthly and annual limits on the number of calls and hours are not violated ($\min\{\cdot\}$ in (2.1b) and (2.1c) is element-wise).

Note on the first day ($d = 1$) of each month m, we set the monthly remaining calls and hours $\mathbf{k}^{\circ}_{m}$ and $\mathbf{h}^{\circ}_{m}$ to their contractual limits $\breve{\mathbf{k}}_{m}$ and $\breve{\mathbf{h}}_{m}$, whereas $\mathbf{X}$ and $\mathbf{Y}$ carry over from the last day of the previous month. To simplify our presentation, we do not show these features in model P.

Problem P is a large-scale intractable integer stochastic dynamic program. Next, we show that even when ignoring the stochasticity, the problem is still strongly NP-hard (Appendix 2.9 provides a mathematical model for the deterministic version of P).

Proposition 1 (NP-Hardness)**.** *The deterministic version of P*

(a) is NP-hard when the number of groups is 2,

(b) is strongly NP-hard when the number of groups is large, and

(c) when the number of groups is large, even the last-day problem is strongly NP-hard.

All proofs are in Appendix 2.10. Our reductions to prove Proposition 1 satisfy the *threshold* policy. Thus, as a corollary, even when we impose a threshold policy on P, it remains strongly NP-Hard. Our numerical experiments indicate only toy sizes of P can be solved in a reasonable time. Therefore, we propose a hierarchical approximation method to solve P. §2.3.1 provides an overview of our approximation. §2.3.2 presents a lower bound for problem P to measure the total error of our approach.

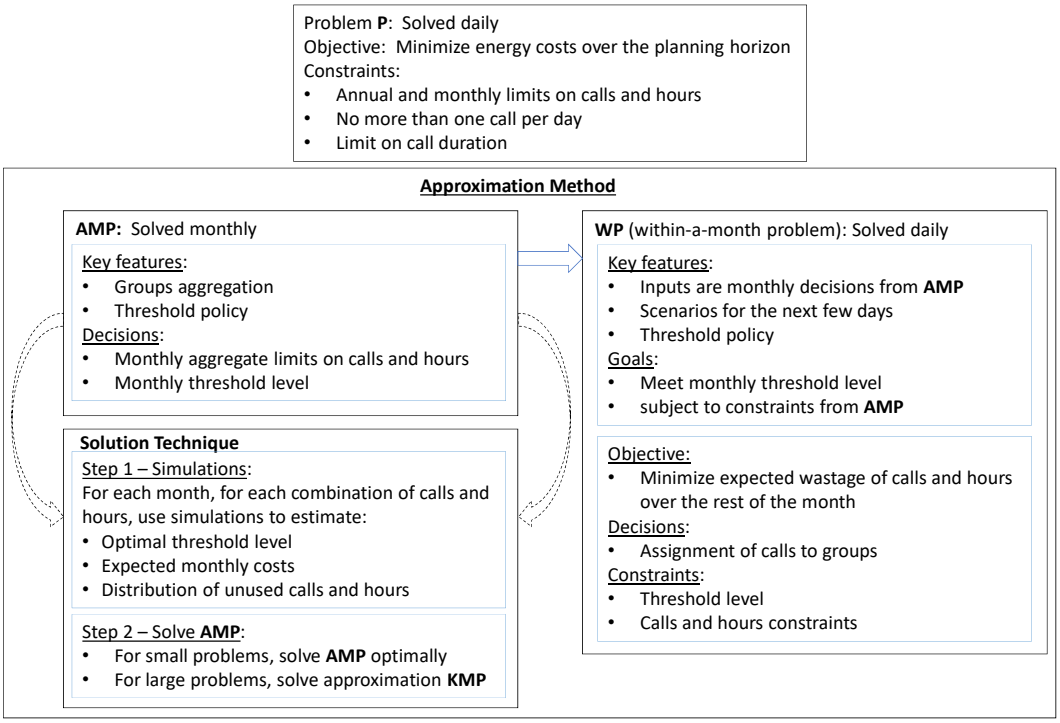

Figure 2.4: Overview of our approximation method

2.3.1 Approximation method

Our approximation method consists of two major components (Figure 2.4). The first component, referred to as the *aggregate monthly problem*, considers the entire planning horizon, and aggregates the groups and the days in each month. The second problem, called the *within-a-month problem*, considers a month as its horizon and determines the calls to be made each day. Decisions made by AMP are constraints for WP. Our approximation is akin to hierarchical approaches to solving production planning problems, which have effectively solved complex practical problems [BHH82, NL92]. At the "master level" (i.e. AMP), we allocate calls and hours to each month, taking into account the monthly and annual contractual limits, and determine the monthly threshold levels. The "operation level" (i.e. WP) determines the call for each group on a daily basis, while incorporating the decisions made by the master problem. In the following, we provide an overview of AMP and WP.

2.3.1.1 Aggregate monthly problem

AMP has the following key elements: (i) We reduce the state space by aggregating the groups and creating a "super group," where the calls and hours available to each group are summed and assigned to this super group; (ii) we reformulate the annual problem in terms of monthly problems; (iii) we employ a threshold policy to compute the monthly costs; that is, we compute a target threshold for each month, and each day, we schedule calls so as to reduce the ECP to this threshold; and (iv) at the beginning of each month, we re-optimize AMP on a rolling-horizon basis. AMP is

$$\text{(AMP)} \qquad V_m(X, Y, \Omega_m) \;=\; \min\{\Phi_m(k, h, \mathcal{Z}, \Omega_m) \;+\; \mathbf{E}[V_{m+1}(X - \tilde{k}, Y - \tilde{h}, \tilde{\Omega}_{m+1})]\}, \tag{2.2a}$$

$$\text{s.t.} \quad \tilde{k} \;\le\; k \;\le\; \min\{X, G\breve{k}_m\}, \tag{2.2b}$$

$$\tilde{h} \;\le\; h \;\le\; \min\{Y, G\breve{h}_m\}, \tag{2.2c}$$

where $\Phi_m(\cdot)$ is the expected cost incurred in month m under forecast information Ω_m, given k and h available calls and hours, respectively, in the month, and while employing threshold $\mathcal{Z}$. The decisions made at the monthly level are k, h, and the threshold $\mathcal{Z}$. These three decision variables are scalars for a given month. Random variables $\tilde{k}$ and $\tilde{h}$ denote the number of calls and hours actually used in the month, determined by the demand realization over the month. The optimal value of AMP is $V_1(GK, GH, \Omega_1)$. We discuss the key features of AMP below.

Groups aggregation. In AMP, state-space aggregation simplifies the problem by reducing the monthly decisions to three scalars. Theorem 2 in §2.4.2 establishes a worst-case error bound for the deterministic version of AMP. Our numerical experiments in §2.6 indicate the error of aggregating groups is small.

Reformulating P in terms of monthly problems. An important aspect of AMP is that it focuses on monthly decisions. We decided to focus on monthly analysis for three reasons.

First, the contractual terms stipulate monthly limits on calls ($\breve{k}_m$) and hours ($\breve{h}_m$). AMP sets limits on the number of monthly calls (k) and hours (h), which may be less than the monthly contractual limits. These decision variables are not present in P. Second, energy-consumption patterns are seasonal and include a significant monthly component. Extended forecasts beyond a few weeks are largely independent of the current state and depend on historical patterns for that time of the year. Third, month is also a suitable unit of analysis for procurement planning purposes. Utilities are concerned with the peak demands in a month, which also justifies employing the monthly threshold levels.

Threshold policy. The industry practice of using a threshold-type policy is motivated by the following advantages. First, on the supply side, utilities commonly need to provide a monthly target plan to their suppliers. Second, a threshold-type policy flattens the demand by reducing the consumption during peak periods. Third, besides the above practical and managerial benefits, a threshold-type policy is also desirable from the view of optimality and computational effectiveness. Given that the cost curve is convex increasing in the height of the energy consumption profile, reducing higher levels of an ECP is clearly preferable.

Overall, our approximation AMP reduces some of the computational burden and has the flexibility to adapt to forecast changes. Aggregating groups facilitates determining the monthly threshold levels and the monthly calls and hours requirements, while considering the needs over the entire planning horizon. The threshold policy, besides being intuitively appealing, simplifies the daily and monthly problems ($\phi(\cdot)$ and $\Phi(\cdot)$). AMP also allows us to incorporate forecast revisions. The daily problem takes into account the latest forecast and, by resolving AMP each month, we incorporate forecast updates at a monthly level.

Next, we provide an overview of our approach to solving AMP (see §2.4 for details).

1. We use simulation to determine the expected cost $\Phi_m(\cdot)$ for each month and each possible value of h, k, and threshold level $\mathcal{Z}$. These simulations also provide an estimate of the distribution of the number of calls and hours that will remain unused at the

end of the month. These inputs are sufficient for solving AMP. Our numerical study indicates only medium instances of AMP can be solved in a reasonable amount of time. Thus, for large instances, we develop an approximation of AMP.

2. We approximate the value function of AMP by ignoring the overflow of unused calls and hours. The resulting approximation of AMP, referred to as KMP, has a nice structure. KMP uses the expected monthly costs $\Phi_m(\cdot)$, obtained from our simulations, and approximates the objective function values of AMP by setting the unused calls and hours ($[k - \tilde{k}]^+$ and $[h - \tilde{h}]^+$) to zero. KMP has the structure of a *multiple-choice knapsack problem*, and its objective function only contains the monthly cost functions $\Phi_m(\cdot)$. Furthermore, our analysis revealed the monthly cost functions $\Phi_m(\cdot)$ are nearly convex. Remarkably, large industrial instances of KMP are solved in a reasonable time. Theorem 3 in §2.4.3 analyzes the properties of KMP.

3. In our simulations, we need to repeatedly solve the daily problems $\phi(\cdot)$ to determine the expected monthly costs $\Phi_m(\cdot)$. Because we are restricting solutions to threshold-type policies, the daily problems minimize the calls and hours needed each day to lower the energy-consumption levels to the given threshold. The daily problems are solved many times in the course of simulation; hence, they must be solved efficiently. We present a polynomial-time procedure for optimally solving these daily problems in §2.4.1.

2.3.1.2 Within-a-month problem

AMP (KMP) determines the threshold level for and the total number of calls and hours that are available in a month. These calls have to be allocated among the G groups to minimize costs in the month, while respecting the contractual limits for each group. Our allocation problem is more difficult than the stochastic bin packing problems [BBD20]. Thus, we design another approximation method, called the *within-a-month problem*, to deploy calls to groups on a daily basis. WP employs a scenario-based rolling-horizon approach, and has the following

key features.

1. We simplify the problem by truncating the horizon. Typically, demand can accurately be forecast for the focal day. For the next few days, demand forecasts have greater uncertainty. Nevertheless, over a horizon consisting of a few days, we can assign likelihoods to alternate realizations. Beyond this small horizon, forecasts resemble what is broadly expected for the month and lack any additional refinements. We assume demand is independent and identically distributed outside a short horizon of a few days. Our scenario-based approach is based on this structure of the forecast evolution.

2. We present a heuristic procedure for solving WP, referred to as HWP, which focuses on minimizing the expected number of hours that remain unused at the end of the month. Under the threshold policy, the number of hours used is approximately proportional to cost reduction. The ability to utilize the hours is constrained by calls. Thus, a poor allocation of calls to groups may result in some groups running out of calls although they may have hours. HWP, which is solved each day, attempts to mitigate such losses.

2.3.2 Lower bound for problem P

We use a lower bound (LB) to measure the total error of our approximation. The LB is obtained by relaxing the constraints on the number of calls while retaining the constraints on hours. We also allow a customer to be called more than once in a day; that is, we ignore the constraints on the continuity of calls. Moreover, we evaluate LB on observed realizations.

Our LB for problem P is

$$LB := \mathbf{E}\left\{\begin{array}{c}\min \\ \sum_{m,d,t}\mathcal{R}_{mdt}\leq G\min\{H,LK\}, \\ \sum_{d,t}\mathcal{R}_{mdt}\leq G\min\{\breve{h}_m,L\breve{k}_m\},\ \forall m, \\ \sum_{t}\mathcal{R}_{mdt}\leq GL,\ \forall m,\ d, \\ \mathcal{R}_{mdt}\in\{0,1,\ldots,G\},\ \forall m,\ d,\ t,\end{array}\left\{\sum_{m,d,t} f_{mdt}\left(\tilde{r}_{mdt}-\mathcal{R}_{mdt}\right)\right\}\right\}.$$

The expectation is over the annual consumption $\{\tilde{r}_{mdt}\}_{m,d,t}$. For a given annual consumption, $\mathcal{R}_{mdt}$ denotes the amount of reduction in consumption during period t in day d of month m. To compute the LB through simulation, for each testing realization, we obtain the optimal $\mathcal{R}_{mdt}$'s that minimize the deterministic optimization problem inside the expectation. For a given realization, the optimal solution is found by assigning $G\min\{H, LK\}$ units of consumption reduction to some periods with the largest cost-reduction over the horizon. This procedure is performed in $\mathcal{O}(GH\mathcal{M}\mathcal{D}_m T)$ for each testing realization (assuming $H \leq LK$). Therefore, the time complexity would be $\mathcal{O}(\mathcal{Q}GH\mathcal{M}\mathcal{D}_m T)$, where $\mathcal{Q}$ denotes the number of testing realizations. Note the LB does not impose the threshold policy, and hence, the errors measured based on this bound incorporate all errors in our approximation, including the errors due to (i) imposing monthly decisions, (ii) restricting solutions to threshold-type policies, (iii) aggregating groups, and (iv) approximately solving WP. Although this bound relaxes two key constraints and is an ex-post optimization, we find the total error of our approximation method, measured against this bound, is small.

2.4 Aggregate Monthly Problem: Solution Methodology

Our dynamic programming problem AMP requires as input, for all possible combinations of calls (k), hours (h), and thresholds ($\mathcal{Z}$), (i) the expected monthly costs $\Phi_m(\cdot)$ and (ii) the distribution of the number of calls and hours used in each month ($\tilde{h}$ and $\tilde{k}$). We estimate these inputs using Monte Carlo simulations. We run one set of simulations for each month, for each combination of monthly allocations of calls and hours, and for each threshold level. Our simulations for each month consist of the following steps:

1. We simulate the ECPs of the days in the month using the historical distribution of the ECPs.

2. For the simulated month, for each combination of monthly allocations of calls and hours, and for each potential threshold level, we compute the daily costs, monthly cost, and

the number of calls and hours that are used in the month.

3. Each day, for each combination of k, h, and $\mathcal{Z}$, we solve a *daily problem*, denoted by DLP, to compute the daily cost. DLP reduces the ECP of the day to the given threshold subject to the availability of calls and hours. By solving DLP, we introduce at most G calls in a day, where each call's duration is at most L. Details are available in §2.4.1.

For each combination of h, k, and $\mathcal{Z}$, a simulated month provides a single estimate of $\Phi_m(\cdot)$, $\tilde{h}$, and $\tilde{k}$. We simulate each month 300 times in our numerical experiments. A computational burden stems from DLP, which determines the day cost $\phi(\cdot)$. We need to solve DLP millions of times; in each iteration, we need to solve DLP for each day of the month, for each combination of h and k, and for each potential threshold level. Next, we present a mathematical model for DLP and an efficient procedure for solving it (§2.4.1), a theoretical and numerical analysis of approximating P by AMP (§2.4.2), and our solution procedures for AMP and its approximation KMP (§2.4.2 and §2.4.3).

2.4.1 Daily problem (DLP) and its solution

For a given ECP, the area above any given threshold is fixed, and hence, we do not have any latitude in the number of hours needed to lower the ECP. Therefore, DLP reduces to a problem that minimizes the number of calls needed to implement the threshold policy.

Let r_{mdt} denote the power-consumption level during period (hour) t of day d in month m. Given r_{mdt}, for $t = 1, \ldots, T$, and a threshold level $\mathcal{Z}$, the total number of hours that have to be chopped off in day d is fixed and computed as $\hat{h}_{md} := \sum_{t=1}^{T} [r_{mdt} - \mathcal{Z}]^+$. Therefore, we need to find the minimum number of calls, denoted by $\hat{k}_{md}$. Let J be a known upper bound on the value of $\hat{k}_{md}$. There are J potential calls, and we aim to determine these calls' starting times and durations. Let binary variable y_{jt} take a value of 1 if the j^{th} potential call starts at period t, and 0 otherwise. Moreover, let binary variable x_{jt} take a value of 1 if the j^{th}

potential call is active during period t, and 0 otherwise. DLP is

$$\text{(DLP)} \qquad \hat{k}_{md} := \min \sum_{j=1}^{J} \sum_{t=1}^{T} y_{jt} \tag{2.3a}$$

$$\text{s.t.} \sum_{j=1}^{J} x_{jt} = [r_{mdt} - \mathcal{Z}]^{+}, \quad \forall t, \tag{2.3b}$$

$$\sum_{t=1}^{T} x_{jt} \leq L, \quad \forall j, \tag{2.3c}$$

$$y_{jt} \geq x_{jt} - x_{j,t-1}, \quad \forall j,\ t \geq 2, \tag{2.3d}$$

$$y_{j1} \geq x_{j1}, \quad \forall j, \tag{2.3e}$$

$$\sum_{t=1}^{T} y_{jt} \leq 1, \quad \forall j, \tag{2.3f}$$

$$x_{jt}, y_{jt} \in \{0, 1\}, \quad \forall j,\ t. \tag{2.3g}$$

The objective function (2.3a) minimizes the total number of calls in day d. Constraint (2.3b) reduces the load to $\mathcal{Z}$.[6] Constraint (2.3c) ensures the duration of each call does not exceed L. Constraints (2.3d)-(2.3f) guarantee the continuity of calls.[7]

A real instance of DLP may have more than 1,000 binary variables. Solving DLP efficiently is essential, given that commercial solvers cannot solve millions of integer programs in a reasonable amount of time. We present a procedure, given in Algorithm 1, that finds an optimal solution for DLP in $\mathcal{O}(JT)$. This procedure finds the minimum number of calls to reduce the load to $\mathcal{Z}$ in day d. The *while*-loop creates calls until the threshold $\mathcal{Z}$ is achieved. The starting time of a call is the earliest time that load level exceeds $\mathcal{Z}$ (line 2), and the call is active until either the length of the call is L or the load level starts reducing (line 3). In line 4, the new call (determined in lines 2-3) is output, and the ECP is updated in line 5.

[6]For ease of presentation, we assume each active call reduces the load by one unit. To relax this assumption, one needs to multiply the left-hand side of constraint (2.3b) by an appropriate coefficient.

[7]If reducing the load to $\mathcal{Z}$ by at most G calls is not possible, the solution of DLP will produce more than G calls. In this case, one could select any G calls from the set of produced calls.

Theorem 1 shows Algorithm 1 finds an optimal solution for DLP in a polynomial time.

Algorithm 1: Finding an optimal solution for DLP

Let $\bar{r}_t := r_{mdt}$, for $t = 1, \ldots, T$, and $\bar{r}_{T+1} := \mathcal{Z}$
while $\exists\ t^s := \min\{t|\ 1 \leq t \leq T,\ \bar{r}_t > \mathcal{Z}\}$ **do**
    Define $t^e := \min\{t^s + L - 1,\ \min\{t|\ t^s \leq t \leq T,\ \bar{r}_t > \bar{r}_{t+1}\}\}$.
    Introduce a call with start and end times t^s and t^e (inclusive).
    Update $\bar{r}_t := \bar{r}_t - 1$, for all $t^s \leq t \leq t^e$.

Theorem 1 (Optimality and Efficiency of the DLP Procedure)**.** *Algorithm 1 finds an optimal solution for DLP in time complexity of* $\mathcal{O}(JT)$*.*

Thus, our procedure optimally solves DLP, and its time complexity is linear in the number of periods T. Theorem 1 is valuable because it enables us to optimally solve millions of instances of DLP in a reasonable amount of time. As discussed earlier, the threshold policy is a heuristic policy, and our numerical experiments verify its high quality; however, theoretically, its worst-case error can be arbitrarily bad. Finally, the outputs of our DLP procedure and the monthly expected costs $\Phi_m(\cdot)$ are inputs for solving AMP (KMP), as we describe next.

2.4.2 Analysis of AMP and a dynamic programming procedure

Our simulation procedure described above provides us with the information needed to solve AMP as a finite-horizon stochastic dynamic program. The computational time complexity of this procedure is $\mathcal{O}(\mathcal{M}G^6\breve{h}_m^2\breve{k}_m^2HK)$, recalling that $\breve{h}_m$ and $\breve{k}_m$ are the maximum number of calls and hours allowed per month per group, and H and K are the annual limits. Our numerical study indicates medium-size instances, with group sizes less than six, can be solved exactly using this dynamic programming procedure. However, large industrial instances (e.g., with 20 groups and annual limits of 100 calls and 180 hours per group) can't be solved in a reasonable amount of time. Therefore, in §2.4.3, we present an approximation of AMP that is suitable for large instances.

Before proceeding, we present an analysis of approximating P by AMP. This step of our approximation introduces an error because groups are aggregated, referred to as *groups-aggregation error*. Below, we present a worst-case bound for this error, for a deterministic special case of our problem.

Theorem 2 (Groups-Aggregation Error)**.** *Let v' and v'' denote the optimal values of the deterministic special cases of P and AMP, respectively, without restricting solutions to threshold-type policies. The gap between these optimal values satisfies:*

$$0 \leq v' - v'' \leq \delta G \left(\mathcal{M}L + \min \left\{ \mathcal{M}L,\ \left(\sum_{m=1}^{\mathcal{M}} \breve{h}_m \right) - H \right\} \right),$$

where δ is an upper bound on the increase in cost if a one-hour call to a group is canceled at any time.

Theorem 2 provides an absolute worst-case error for the groups aggregation. We measured this error using various sets of practical parameters. The relative value of this worst-case error ranges roughly from 30% to 55% (with respect to saving). However, our numerical experiments in §2.6 show the amount of error introduced by aggregating groups is considerably smaller.

Approximating P by AMP introduces another error because of restricting the solutions to the threshold-type policies; that is, each day, we schedule calls so as to reduce the ECP to a target threshold (recall a third type of error is introduced due to imposing monthly decisions). Our numerical results in §2.6 indicate the relative error is on average less than 7.76%.

Thus far, we have provided an analysis of the errors of approximating P by AMP, and argued a dynamic programming algorithm can solve medium-size instances of AMP. Next, we present an approximation of AMP, which enables us to solve large instances in a reasonable amount of time.

2.4.3 Multiple-choice knapsack approximation to AMP: KMP

At the beginning of each month m, AMP allocates calls (k_m^*) and hours (h_m^*) to the month and sets a threshold level ($\mathcal{Z}_m^*$). These decision variables determine the expected monthly costs $\Phi_m(\cdot)$ and the distributions of the unused calls and hours from the allocation at the end of the month. The resources (calls and hours) that are allocated to a month are optimized by considering the expected future value of these resources. Calls and hours available beyond month m depend on k_m^*, h_m^*, and the unused resources left at the end of month m. One source of complexity in AMP is these unused calls and hours. In our approximation KMP, we assume these random variables are zero, which significantly reduces the computational burden and allows us to solve industrial-size instances of KMP in a reasonable time.

Before examining properties of KMP, we discuss the reason we chose this approximation. We conjectured that the saving derived under AMP and KMP are likely to be very similar based on the following intuitive arguments. First, because AMP accounts for unused resources, at the beginning of the horizon, its key decision variable is the threshold level ($\mathcal{Z}_m^*$). AMP can be liberal with resource allocation, recognizing the threshold level will dictate the usage in the focal month and the overflow into the next month. On the other hand, KMP has to carefully match the threshold level and the calls and hours allocated to a month. In KMP, the only decision variables are calls (k) and hours (h) allocated to a month. These decisions determine a threshold that leads to the minimum cost for the allocated resources. Thus, we conjectured that in the early parts of the horizon, KMP and AMP may differ in their resource allocations but will set similar threshold levels, and hence, they will use similar amount of resources. Towards the end of the horizon, we expected AMP and KMP to also match on the resources allocated. Second, by assuming that in KMP the allocated resources (k_m^* and h_m^*) are consumed in the month, we are introducing an error in the expected value of using these resources. Interestingly, this error applies to all possible values of resources, and hence, although the expected costs are incorrectly computed in KMP, to the extent that the same error is made across all combinations of calls and hours, the savings obtained in each

month by AMP and KMP may be the same. Third, although we ignore the unused resources when solving KMP, we incorporate any unused resources when we re-solve the problem every month. Our above conjectures were confirmed by our numerical experiments on medium instances, which indicated the difference between the optimal values of AMP and KMP is less than 8.33% and, on average, 5.14% (see §2.6).

We explore the structure of KMP. Because we assume all allocated resources are consumed in any month, the resulting optimization problem can be represented as the problem of selecting a policy among the set of policies for each month while being mindful of the yearly allocation of calls and hours. This problem can be formulated as a two-dimensional multiple-choice knapsack problem. We show it can be solved effectively due to its special cost function.

We present a mathematical model for KMP. Recall we use simulations to estimate the expected monthly costs $\Phi_m(\cdot)$, for all m, k, h, and $\mathcal{Z}$. If we ignore the overflows from one month to the next, for each combination of h and k, we can pick a threshold level $\mathcal{Z}$ that minimizes the monthly cost. As a result, for ease of exposition, in this section, we denote the expected monthly costs by $\Phi_m(k, h)$. For each month m, and for some values of k and h, let $\mathbf{p}_{m,i}$ denote the triple $(k, h, \Phi_m(k, h))$ and let $p_{m,i,j}$ denote the j^{th} element of this triple. Each month has $N_m = (G\breve{k}_m + 1)(G\breve{h}_m + 1)$ triples. Furthermore, let $\mathbb{P}_m := \{\mathbf{p}_{m,1}, \ldots, \mathbf{p}_{m,N_m}\}$. KMP is a two-dimensional multiple-choice knapsack, where the points $\mathbf{p}_{m,i}$ are analogous to the items in a knapsack problem that are partitioned into $\mathcal{M}$ subsets $\mathbb{P}_1, \ldots, \mathbb{P}_{\mathcal{M}}$, and exactly one item must be selected from each subset. Selecting item $\mathbf{p}_{m,i}$ utilizes the two resources (annual calls and hours), and the amounts of resource utilization are given by the first two elements of the vector $\mathbf{p}_{m,i}$. The third element of $\mathbf{p}_{m,i}$ corresponds to the objective coefficient in the knapsack problem. KMP is formulated as

$$\text{(KMP)} \quad \min \sum_{m=1}^{\mathcal{M}} \sum_{i=1}^{N_m} p_{m,i,3} \nu_{m,i} \tag{2.4a}$$

$$\text{s.t.} \quad \sum_{m=1}^{\mathcal{M}} \sum_{i=1}^{N_m} p_{m,i,1} \nu_{m,i} \leq GK, \tag{2.4b}$$

$$\sum_{m=1}^{\mathcal{M}} \sum_{i=1}^{N_m} p_{m,i,2} \nu_{m,i} \leq GH, \tag{2.4c}$$

$$\sum_{i=1}^{N_m} \nu_{m,i} = 1, \quad \forall m, \tag{2.4d}$$

$$\nu_{m,i} \in \{0,1\}, \quad \forall m, \ i. \tag{2.4e}$$

Constraint (2.4d) ensures exactly one point is selected for each month from the set $\mathbb{P}_m$; that is, we pick one (k,h) pair for each month, satisfying $0 \leq k \leq G\breve{k}_m$ and $0 \leq h \leq G\breve{h}_m$. Constraints (2.4b) and (2.4c) enforce the annual calls and hours constraints, respectively, recalling the first and second entries of the selected points represent the number of calls and hours. The objective function minimizes the sum of the third entries of the selected points, representing the total expected cost over the contract horizon.

The constraints and objective function of KMP have desirable features. We first explore the constraints. In addition to the binary constraints (2.4e), KMP has $\mathcal{M}$ generalized upper-bound constraints (2.4d) and two knapsack constraints (2.4b)-(2.4c). In the LP relaxation of KMP, denoted by KMP°, the binary constraints are redundant and can be dropped, leaving $\mathcal{M}+2$ constraints. Hence, in KMP°, at most $\mathcal{M}+2$ non-zero variables exist that must sum up to $\mathcal{M}$. This observation implies at least $\mathcal{M}-2$ decision variables will be 1 and at most four decision variables will be fractional.

The objective function of KMP consists of the expected costs $\Phi_m(k,h)$. The value of $\Phi_m(k,h)$ is monotone non-decreasing in k and h. Although $\Phi_m(h,k)$ is not fully convex, our numerical analysis indicate it is approximately convex. For example, the left graph in Figure 2.5 depicts the points in $\mathbb{P}_m$ for a small instance with $G=5$, $\breve{k}_m=2$, and $\breve{h}_m=5$. We have connected these points to provide a visualization of the cost function (this step is akin to assuming calls and hours take continuous values). Note the values have been scaled to [0,1]

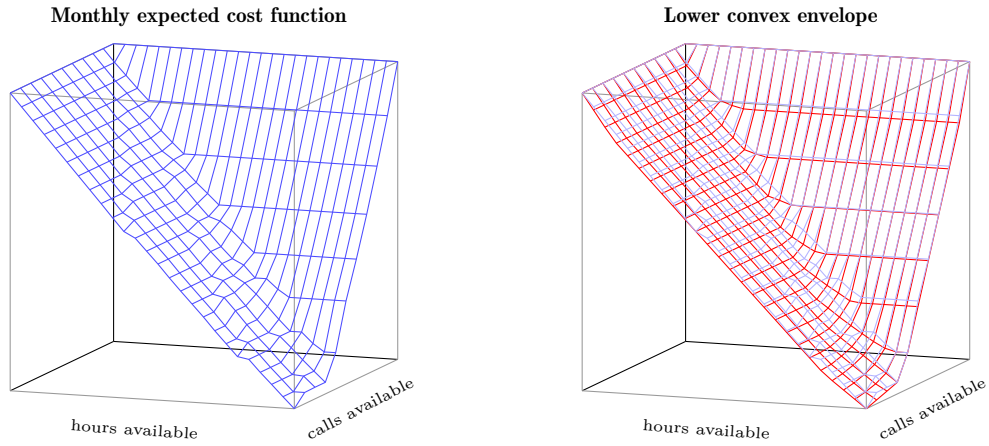

Figure 2.5: A monthly expected cost function (a) and its lower convex envelope (b)

to provide an improved visual illustration.

We next formalize the relationship between KMP and its LP relaxation KMP°. We show that under some conditions, KMP° solves KMP. Let $\mathbf{p}_m^\circ := \sum_{i=1}^{N_m} \mathbf{p}_{m,i}\nu_{m,i}$ denote the selected point for month m in the optimal solution of KMP°. Recall the elements of $\mathbf{p}_m^\circ$ are respectively the number of calls ($p_{m,1}^\circ$), hours ($p_{m,2}^\circ$), and cost ($p_{m,3}^\circ$) for month m. The objective function forces $\mathbf{p}_m^\circ$ to lie on the *lower envelope* of the convex hull generated by the set $\mathbb{P}_m$. Let $conv(\mathbb{P}_m)$ denote the convex hull of $\mathbb{P}_m$. The lower envelope of $conv(\mathbb{P}_m)$ is defined as

$$\mathcal{E}_m : [0, G\breve{k}_m] \times [0, G\breve{h}_m] \to \mathbb{R},\ \mathcal{E}_m(p_{m,1}, p_{m,2}) = \min\{p_{m,3} \mid \mathbf{p}_m \in conv(\mathbb{P}_m)\}.$$

The right graph in Figure 2.5 depicts the lower convex envelope. The points in $\mathbb{P}_m$ are also plotted in a lighter tone. Observe the lower convex envelope is very close to $\mathbb{P}_m$. We consistently observed this in all of our numerical experiments, implying KMP° provides a very good bound for KMP. In fact, under some reasonable assumptions, we can find an optimal solution for KMP by solving KMP°, as presented below.

Theorem 3 (Properties of KMP°)**.** *The optimal solution of KMP° satisfies the following*

properties.

(a) In an optimal solution of KMP°, for all m, $\mathbf{p}^{\circ}_m$ is on the lower convex envelope $\mathcal{E}_m$.

(b) If for all m, the lower envelope $\mathcal{E}_m$ is linear over the unit square $[k, k+1] \times [h, h+1]$, for all $k \in \{0, 1, \ldots, G\breve{k}_m - 1\}$ and $h \in \{0, 1, \ldots, G\breve{h}_m - 1\}$, then KMP° has an optimal solution with integral numbers of calls and hours for all months.

(c) KMP° has an optimal solution that is also optimal for KMP if for all m, the following conditions are satisfied: (i) The lower convex envelope $\mathcal{E}_m$ is linear over the unit square $[k, k+1] \times [h, h+1]$, for all $k \in \{0, 1, \ldots, G\breve{k}_m - 1\}$ and $h \in \{0, 1, \ldots, G\breve{h}_m - 1\}$; and (ii) $\mathbf{p}_{m,i}$ is on the lower envelope $\mathcal{E}_m$, for all i.

In summary, we showed KMP has several nice structural properties that make it easier to solve than AMP. In our numerical experiments on a large practical-size instance with 25 groups, $\breve{k}_m = 10$, $\breve{h}_m = 20$, $K = 100$, and $H = 180$, an instance of KMP has, on average, 28 rows, 817,401 columns, 817,382 integer variables, and 3,264,671 nonzeros, and is solved in 15.52 seconds, on average, on a personal computer. For this instance, computing the expected monthly costs through simulation (to construct KMP instances), using 300 realizations, took 53,785.49 seconds on our personal computer. Note computing these monthly costs can be performed *offline* and *on parallel machines*, which can significantly reduce its CPU time. Furthermore, we show, on medium-size instances, the error gap in approximating AMP by KMP is small (less than 8.33% and, on average, 5.14%; see §2.6). By solving AMP (KMP) at the beginning of each month, we identify the number of calls and hours to be used during the month.[8]

[8]We also investigated the impact of partially enumerating these points on the quality of our solution. On the instances of Table 2.1, enumerating only 10% of the points (randomly selected) increased the error of KMP+HWP by less than 1%. Thus, we conclude our approximation's high performance is preserved if only a small portion of the points are enumerated. We provide an intuitive explanation using Figure 2.5, which shows monthly expected cost and its lower convex envelope. By eliminating some of the points, the lower convex envelope will not change much. For large instances, the number of points is significantly large (e.g. 10,000 points), yet having only a small portion of these points suffices to achieve a good approximation of the envelope. Therefore, especially for large instances, enumerating only a small portion of the points is a reasonable approximation.

2.5 Within-a-Month Problem: Solution Approach

AMP (KMP) allocates aggregate resources (total number of calls k_m^* and hours h_m^*) to each month. These aggregate resources must be disaggregated across the G groups on a daily basis within each month. In other words, each day, we need to assign calls to specific groups so as to reduce the ECPs to the threshold level $\mathcal{Z}_m^*$ obtained from AMP. Our WP is

$$\begin{aligned}
\text{(WP)} \qquad w_{m,d}(\mathbf{X}, \mathbf{Y}, \mathbf{k}_m^\circ, \mathbf{h}_m^\circ, \Omega_{m,d}) \;=\; & \min\{\breve{\phi}_{m,d}(\mathbf{k}, \mathbf{h}, \mathcal{Z}_m^*, \Omega_{m,d}) \\
& + \; \mathbf{E}[w_{m,d+1}(\mathbf{X}-\mathbf{k}, \mathbf{Y}-\mathbf{h}, \mathbf{k}_m^\circ - \mathbf{k}, \mathbf{h}_m^\circ - \mathbf{h}, \tilde{\Omega}_{m,d+1})]\}, \qquad (2.5a) \\
\text{s.t.} \quad & \text{Constraints (2.1b)-(2.1c), and} \\
& \mathbf{k}^\top \mathbf{1} \;\le\; k_m^\triangle \;:=\; k_m^* - \left(G\breve{k}_m - \mathbf{k}_m^{\circ\top}\mathbf{1}\right), \qquad (2.5b) \\
& \mathbf{h}^\top \mathbf{1} \;\le\; h_m^\triangle \;:=\; h_m^* - \left(G\breve{h}_m - \mathbf{h}_m^{\circ\top}\mathbf{1}\right), \qquad (2.5c)
\end{aligned}$$

recalling that $\mathbf{X}$ and $\mathbf{Y}$ (respectively, $\mathbf{k}_m^\circ$ and $\mathbf{h}_m^\circ$) denote the calls and hours available to each group based on the annual (respectively, monthly) contractual limits. The daily cost $\breve{\phi}_{m,d}(\cdot)$ is the cost of supplying power after the ECP is truncated to the threshold. A solution of WP must satisfy constraints (2.1b)-(2.1c) to ensure the monthly and annual contractual limits for each group are not violated. In addition, constraints (2.5b)-(2.5c) ensure the aggregate monthly limits k_m^* and h_m^* obtained from AMP (KMP) are adhered to. Note $k_m^\triangle$ and $h_m^\triangle$ denote the number of calls and hours left based on the limits k_m^* and h_m^* set by AMP (KMP).

Recall Algorithm 1 solves DLP and determines the calls needed to reduce the load to the given threshold level. WP assigns these calls to the groups. WP is different from P because (i) in WP, the horizon is limited to a month, and (ii) each day, groups have to be called to reduce the ECP to the given threshold level. Overall, WP is similar in structure to P, and it is another difficult stochastic dynamic program. We use a reduction similar to that in Proposition 1 to show WP too is strongly NP-hard.

Proposition 2 (WP is NP-hard)**.** *The deterministic special case of WP is strongly NP-hard,*

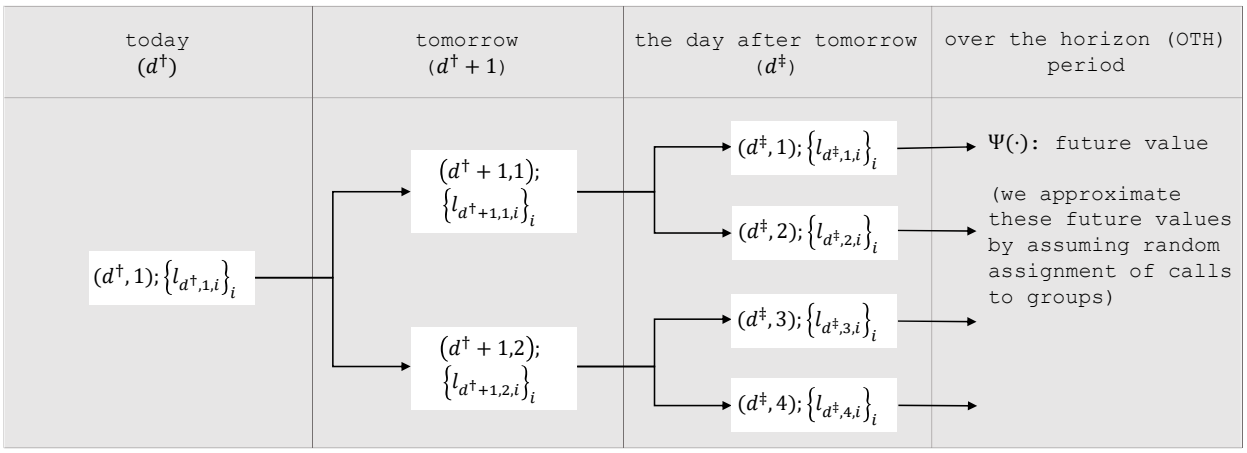

Figure 2.6: Illustration of our scenario-based disaggregation model

implying WP is strongly NP-hard.

Moreover, our numerical experiments indicate only toy-size instances of WP can be solved optimally using commercial solvers (e.g., using the stochastic dynamic programming approach). Thus, we design a heuristic approach for solving WP, referred to as HWP.

Energy demand can be forecasted with some accuracy over a short horizon of a few days. Beyond this horizon, forecasts resemble what would be expected for the month and day of the week. Motivated by this observation, we present our heuristic HWP to solve WP, which uses a scenario-based approach. In our scenario-based approach (Figure 2.6), we assume the demand for the focal day $d^\dagger$ is known with certainty, and the forecast for the next few days $d^\dagger+1,\ldots,d^\ddagger$ consists of discrete alternative demand realizations or scenarios. Beyond day $d^\ddagger$, we assume the demand is independent and identically distributed. We refer to the part of the horizon that is beyond $d^\ddagger$ as *over-the-horizon* (OTH) period.

Let (d,s) denote the s^{th} potential demand realization on day $d\in\{d^\dagger,\ldots,d^\ddagger\}$. A specific ECP is associated with each scenario (d,s). Thus, associated with each scenario is a set of calls, denoted by $\{l_{d,s,i}\}_i$, that are obtained by solving DLP for the corresponding ECP. We need to assign specific groups to these calls. Let ρ denote a path through the scenarios; for example, in Figure 2.6, scenarios $(d^\dagger,1)$, $(d^\dagger+1,2)$, and $(d^\ddagger,3)$ constitute a path. All paths

start on day $d^\dagger$ and terminate on day $d^\ddagger$. Let π_ρ denote the probability of path ρ.

We first provide an intuition for the design of HWP. Recall WP aims to reduce the ECPs to the given threshold $\mathcal{Z}_m^*$ using the given calls k_m^* and hours h_m^*. A disaggregation method is optimal if either of the following two (sufficient but not necessary) conditions hold: (i) All ECPs are below the threshold when we exhaust the hours h_m^*; or (ii) at the end of the month, all ECPs are below the threshold. A disaggregation method may not be optimal if the allocations result in some groups that have calls, based on their monthly and annual constraints, but do not have hours, and other groups that have hours but have reached their limits on calls, while some ECPs are above the threshold. Thus, as a heuristic approach, we seek to maximally cover the calls whereas minimizing the hours that go unused. We model HWP as

$$\text{(HWP)} \qquad \min \quad \sum_{\rho} \pi_\rho \left(\mathcal{U}_\rho + \Psi_{m,d}\left(\mathbf{X}-\mathbf{k}_\rho, \mathbf{Y}-\mathbf{h}_\rho, \mathbf{k}^\circ-\mathbf{k}_\rho, \mathbf{h}^\circ-\mathbf{h}_\rho\right)\right) \tag{2.6a}$$

$$\text{s.t.} \quad \mathcal{U}_\rho = \sum_{(d,s)\in\rho} \sum_{i} l_{d,s,i} \left(1 - \sum_{g} \alpha_{d,s,i,g}\right), \quad \forall \rho, \tag{2.6b}$$

$$\sum_{(d,s)\in\rho} \sum_{i} \lambda_{d,s,i,g} = k_{g,\rho} \leq \min\left\{X_g, k_g^\circ\right\}, \quad \forall \rho,\ g, \tag{2.6c}$$

$$\sum_{(d,s)\in\rho} \sum_{i} l_{d,s,i}\alpha_{d,s,i,g} = h_{g,\rho} \leq \min\left\{Y_g, h_g^\circ\right\}, \quad \forall \rho,\ g, \tag{2.6d}$$

$$\sum_{g} \sum_{(d,s)\in\rho} \sum_{i} \lambda_{d,s,i,g} \leq k_m^\triangle, \quad \forall \rho, \tag{2.6e}$$

$$\sum_{g} \sum_{(d,s)\in\rho} \sum_{i} l_{d,s,i}\alpha_{d,s,i,g} \leq h_m^\triangle, \quad \forall \rho, \tag{2.6f}$$

$$\sum_{i} \lambda_{d,s,i,g} \leq 1, \quad \forall\ (d,s),\ g, \tag{2.6g}$$

$$\sum_{g} \lambda_{d,s,i,g} \leq 1, \quad \forall\ (d,s),\ i, \tag{2.6h}$$

$$0 \leq \alpha_{d,s,i,g} \leq \lambda_{d,s,i,g}, \quad \forall\ (d,s),\ g,\ i, \tag{2.6i}$$

$$\lambda_{d,s,i,g},\ \in \{0,1\}, \quad \forall\ (d,s),\ g,\ i. \tag{2.6j}$$

The objective function (2.6a) minimizes the expected uncovered hours in the given calls and the expected number of hours that will remain unused in the OTH period, respectively denoted by $\mathcal{U}_\rho$ and $\Psi_{m,d}(\cdot)$. We will shortly describe our approach for estimating $\Psi_{m,d}(\cdot)$. Equation (2.6b) computes the total unassigned hours $\mathcal{U}_\rho$ for each path, where $l_{d,s,i}$ denotes the length of the i^{th} call and variable $\alpha_{d,s,i,g}$ denotes the fraction of call i assigned to group g in scenario (d, s). The continuous variables $\alpha_{d,s,i,g}$ allow for partial assignments of calls; for example, if a group has one call and two hours available, we may assign half of a call with length 4 to this group ($\alpha_{d,s,i,g} = 0.5$). If each of the calls is fully assigned, $\sum_g \alpha_{d,s,i,g}$ will be 1 for each call and $\mathcal{U}_\rho$ will be 0. We also define binary variable $\lambda_{d,s,i,g}$ as 1 if (a non-zero fraction of) call i is assigned to group g in scenario (d, s), and 0 otherwise. In essence, $\lambda_{d,s,i,g}$ accounts for calls used by group g. Constraints (2.6c) and (2.6d) ensure, in each path, the calls and hours assigned to each group are constrained by the annual and monthly contractual limits (recall vectors $\mathbf{X}$, $\mathbf{Y}$, $\mathbf{k}^\circ$, and $\mathbf{h}^\circ$ represent the number of calls and hours available with respect to the annual and monthly contractual limits). Constraints (2.6e) and (2.6f) ensure the calls and hours assigned in each path do not exceed the monthly limits (k_m^* and h_m^*) set by AMP (KMP). Due to constraints (2.6g) and (2.6h), a group may receive at most one call per day, and a call is assigned to at most one group.

The core challenges encounterd while computing the expected loss of hours $\Psi_{m,d}(\cdot)$ in the OTH period are the same as those in WP. Therefore, we use the following simplifications to estimate the loss function:

1. We assume demand (realization of ECPs) is i.i.d. in the OTH period. This assumption implies a stochastic process that generates a set of calls $\{l_i\}_i$ each day.

2. To cover these calls, we rotate through the groups one at a time and randomly assign a call. This procedure is equivalent to shuffling the calls and dealing them to the groups. The rotational schedule is maintained across days. We continue to assign calls until we reach the limits set by AMP (KMP), either k_m^* or h_m^*, and ignore all other constraints.

All calls that are assigned but cannot be implemented due to other constraints are considered lost. We refer to this procedure as the *dealing algorithm* (DA).

3. We assume the loss function $\Psi_{m,d}(\cdot)$ is additive in the groups and let

$$\Psi_{m,d}\left(\mathbf{X},\mathbf{Y},\mathbf{k}^{\circ},\mathbf{h}^{\circ}\right)=\sum_{g}\psi_{m,d,g}\left(X_g,Y_g,k_g^{\circ},h_g^{\circ},\frac{k_m^{\triangle}}{G},\frac{h_m^{\triangle}}{G}\right), \tag{2.7}$$

where function $\psi_{m,d,g}(\cdot)$ computes the expected hours lost by group g. Recall $k_m^{\triangle}$ and $h_m^{\triangle}$ denote the number of calls and hours left based on the limits set by AMP (KMP). Note DA assigns $\frac{k_m^{\triangle}}{G}$ calls to each group, unless the total hours assigned to the group exceeds $\frac{h_m^{\triangle}}{G}$. This procedure ensures the calls and hours allocated by AMP (KMP) are fully deployed. The availability of calls and hours for each group based on their respective annual and monthly constraints are captured in the first four arguments of $\psi_{m,d,g}(\cdot)$.

These simplifications enable us to solve HWP. We next provide a rationale for our approach and then elaborate on the solution procedure. Observe if the limits set by AMP (KMP) are achieved without any hours going unused, DA achieves the same cost as the aggregate solution. Thus, we incur a loss if some groups are unable to accept a call before either of the AMP (KMP) limits are reached. In the following theorem, we show if calls and hours available to a group are in balance, the relative loss in hours due to DA asymptotically approaches zero.

Theorem 4 (Asymptotic Optimality of DA)**.** *Assume calls are assigned to groups according to DA, and have a mean $\bar{l}$ and a standard deviation σ_l. Further, assume*

(a) the monthly limits by AMP (KMP) satisfy $h_m^* = \bar{l}\, k_m^*$,

(b) each group has $\frac{h_m^*}{G}$ hours and $\frac{k_m^*}{G}$ calls available (assuming these quantities are integral), and

(c) $\lim_{n\to\infty}\frac{\sum_{i=1}^{n} l_i - n\bar{l}}{\sigma_l\sqrt{n}} \xrightarrow{\mathscr{D}} \mathcal{N}(0,1)$.

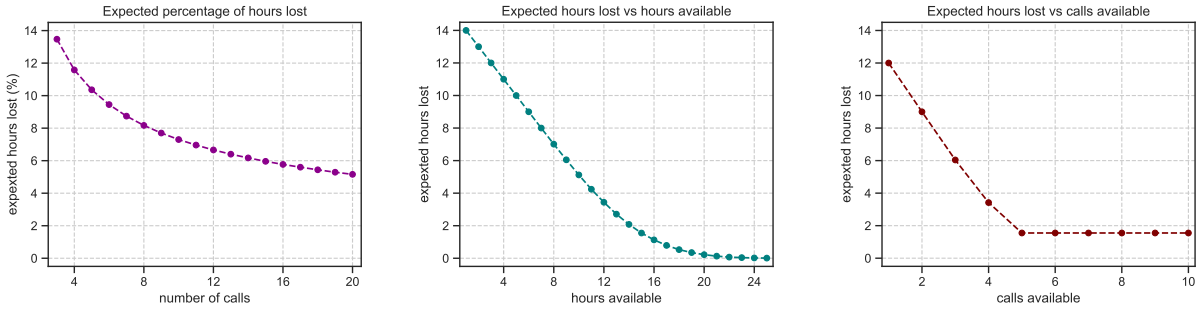

(a) Hours lost (asymptotics) (b) Hours lost vs. hours available (c) Hours lost vs. calls available

Figure 2.7: Expected loss in hours (assume $\bar{l} = 3$, $\sigma_l = 1.73$, and a call's maximum duration L is 6 hours). In Figure (b), we assume calls available are equal to 5. Similarly, in Figure (c), we assume hours available are equal to 15.

Consider a regime in which the length of the OTH period, the number of calls k_m^*, and the number of hours h_m^* grow proportionally. In this regime, let n denote the calls and $n\bar{l}$ denote the hours available per group and also the AMP (KMP) limits per group. Then, the expected percentage of hours lost (computed based on the total hours available) converges to

$$\lim_{n\to\infty} \frac{\beta\sigma_l}{\sqrt{n}\,\bar{l}} + \frac{\epsilon}{n},$$

where β is a constant coefficient and $\epsilon \in [0, 1)$. Thus, DA is asymptotically optimal.

Theorem 4 establishes that if calls and hours are in balance, the percentage error of DA asymptotically approaches zero as the problem size grows. By balance, we mean each group has the number of calls (n) required by AMP (KMP) and the hours available are $n\bar{l}$. Figure 2.7a illustrates that even for a small number of calls n, the expected percentage of hours lost monotonically decreases. In this figure, for each n on the x-axis, the available hours are equal to $n\bar{l}$. Figures 2.7b and 2.7c depict function $\psi(\cdot)$ and demonstrate the impact of an imbalance in calls and hours available to a group. Figure 2.7b shows the change in hours lost as a function of the number of hours available. In this graph, the number of calls are fixed to five. Figure 2.7c shows the impact of the number of calls available. Here, the hours are kept constant at 15. The expected hours lost decline sharply when the number of calls available

equals five. We make the following observations.

1. If calls and hours are in balance, meaning each group has the calls needed by AMP (n) and the available hours are close to $n\bar{l}$, disaggregating the calls and hours such that we realize cost reductions that are close to those obtained in the aggregate model is possible. A simple algorithm such as DA can achieve such a disaggregation.

2. HWP is likely to allocate calls in a balanced manner because $\psi(\cdot)$ heavily penalizes imbalances.

3. When the calls and hours available are in balance, DA is nearly optimal. Hence, DA is likely to provide good estimates of the expected hours that are likely to be lost.

4. DA decomposes the loss function $\Psi(\cdot)$ by groups $\psi(\cdot)$, which considerably reduces the state space (for estimating the loss function) from $\mathcal{O}(\breve{h}_m^G \breve{k}_m^G)$ to $\mathcal{O}(\breve{h}_m \breve{k}_m)$.

5. A balanced calls allocation is also a fair allocation policy and managerially appealing.

Our approach to solving HWP consists of the following two steps. First, we enumerate $\psi(\cdot)$, for all values of available calls and hours, using simulation. This step is performed in $\mathcal{O}(\breve{h}_m \breve{k}_m)$ iterations (note $\psi(\cdot)$ is identical across groups). Second, we reformulate HWP as an integer linear program by introducing binary variables that select one value of $\psi(\cdot)$ for each group.

We construct and solve HWP at the beginning of a day if calls are to be assigned to some groups, that is, if the load exceeds the threshold. In our numerical experiments on a large practical-size instance with 25 groups, $\breve{k}_m = 10$, $\breve{h}_m = 20$, $K = 100$, and $H = 180$, an instance of HWP has, on average, 5,156 rows, 22,020 columns, 17,709 integer variables, and 67,789 nonzeros, and the average CPU time for finding a solution with a 1% optimality gap is 8.01 seconds on a personal computer.

The solution of HWP is the final solution of our approximation method. Intuitively, the total cost associated with this final solution must be worse than the objective value of

AMP (KMP), because HWP disaggregates the solution of AMP (KMP), which may result in some calls and/or hours being wasted (e.g., if a recipient group for a call has reached its annual/monthly limit on either calls or hours). However, HWP has a *forward-looking* feature that enables skipping shorter calls in anticipation of longer (more valuable) calls in the future. If the consumption in a month is unusually higher than expected, obviously, the allocated resources by AMP (KMP) will not be sufficient to reduce the load to the given threshold throughout the month, and the load will inevitably exceed the threshold on some days. In this case, HWP properly rations the available number of calls and does not waste them on short calls. This feature becomes more advantageous when the limit on the number of calls becomes more restrictive. Our numerical experiments on medium-size instances show the objective value of the HWP solution is always better than the objective value of the aggregate solution of AMP (KMP). On our medium instances, HWP improves the solution of AMP by 0.24% on average, and as much as 2.81%. Thus, the forward-looking feature of HWP significantly improves the solution of AMP (KMP), and the improvement is larger than the groups-aggregation error. We present our numerical study below.

2.6 Implementation and Computational Experiments

We conduct extensive numerical experiments to test the performance of our approximation methods and analyze the impact of different problem parameters. We consider 25 instances that are generated by varying the parameters of our problem (G, L, $\breve{k}_m$, $\breve{h}_m$, K, and H). In the following, we will discuss our parameter settings as well as the suggested approach for simulating the training and testing realizations. We summarize our key findings as follows: First, the error (in the objective value) of approximating AMP with KMP ranges between 3.53% and 8.33%, with an average of 5.14%. This is the error in the objective value of the aggregate monthly problem. Second, the total error ranges between 9.82% and 14.33%, with an average of 12.03%. To compute the total error, we compared the solution obtained from

our approach (KMP+HWP) with our lower bound (§2.3.2). This error is remarkably small for such a complex stochastic dynamic program. Further, the error is relative to a lower bound that relaxes most of the constraints. Third, we find our approach (KMP+HWP) allocates calls and hours in a balanced manner, which is consistent with our motivation for designing HWP. The details are presented below.

Table 2.1: Performance analysis of our approximation approaches

						%error vs AMP			%error vs LB				coefficient of variation (%)				
G	L	$\check{k}_m$	$\check{h}_m$	K	H	KMP	AMP + HWP	KMP + HWP	AMP	KMP	AMP + HWP	KMP + HWP	LB	AMP	KMP	AMP + HWP	KMP + HWP
(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)	(i)	(j)	(k)	(l)	(m)	(n)	(o)	(p)	(q)	(r)
2	3	5	10	20	40	4.42	−0.04	4.26	8.15	12.21	8.11	12.06	3.82	4.04	3.57	4.05	3.55
3	3	5	10	20	40	5.26	−0.05	5.06	8.15	12.99	8.11	12.81	3.75	4.38	3.90	4.36	3.87
4	3	5	10	20	40	5.60	−0.14	5.28	7.53	12.71	7.40	12.41	3.57	3.99	3.80	4.00	3.79
5	3	5	10	20	40	4.68	−0.16	4.26	7.20	11.55	7.05	11.16	3.56	4.17	3.75	4.15	3.68
6	3	5	10	20	40	4.23	−0.19	3.77	6.69	10.64	6.52	10.21	3.74	3.75	3.69	3.77	3.68
4	2	5	10	20	40	3.53	−2.81	−0.36	13.30	16.36	10.86	12.98	3.47	4.79	4.10	4.64	3.74
4	3	5	10	20	40	4.31	−0.10	4.00	7.98	11.94	7.88	11.66	3.90	4.47	4.05	4.50	3.99
4	4	5	10	20	40	5.76	−0.15	5.29	7.25	12.59	7.11	12.15	3.56	3.94	3.86	3.97	3.82
4	5	5	10	20	40	4.54	−0.26	3.94	7.01	11.24	6.77	10.67	3.71	4.25	3.77	4.31	3.74
4	6	5	10	20	40	3.67	−0.21	2.98	7.05	10.46	6.85	9.82	3.83	4.36	3.92	4.43	3.90
4	3	3	6	12	24	5.28	−0.24	4.84	8.03	12.88	7.80	12.48	4.12	4.55	3.92	4.53	3.92
4	3	4	8	16	32	5.55	−0.09	5.21	7.74	12.86	7.66	12.55	3.96	4.05	4.00	4.08	4.00
4	3	5	10	20	40	4.66	−0.11	4.35	7.68	11.98	7.58	11.69	3.63	4.31	4.24	4.34	4.18
4	3	6	12	24	48	4.61	−0.12	4.32	7.52	11.79	7.41	11.52	3.45	4.06	3.80	4.06	3.77
4	3	7	14	28	56	3.96	−0.08	3.72	7.46	11.12	7.38	10.91	3.71	4.11	4.09	4.09	4.07
4	3	5	8	20	32	5.64	−0.17	5.19	7.68	12.89	7.53	12.47	3.97	4.47	4.04	4.52	3.99
4	3	5	9	20	36	4.58	−0.21	4.17	7.63	11.86	7.43	11.48	3.80	4.46	4.03	4.47	4.00
4	3	5	10	20	40	5.48	−0.08	5.14	7.47	12.54	7.39	12.22	3.70	4.36	3.72	4.34	3.69
4	3	5	11	20	44	5.36	−0.12	5.09	7.85	12.79	7.74	12.54	3.66	4.29	3.87	4.29	3.86
4	3	5	12	20	48	4.97	−0.07	4.70	7.51	12.11	7.44	11.86	3.22	3.71	3.68	3.68	3.66
4	3	5	10	10	20	7.26	−0.12	6.88	6.33	13.13	6.21	12.77	4.98	5.20	4.17	5.17	4.16
4	3	5	10	15	30	8.33	−0.10	7.90	6.99	14.74	6.89	14.33	4.33	4.14	3.75	4.13	3.75
4	3	5	10	20	40	4.95	−0.11	4.62	7.25	11.84	7.15	11.53	3.55	3.79	3.90	3.79	3.87
4	3	5	10	25	50	6.28	−0.10	5.90	8.23	13.99	8.14	13.64	3.43	3.80	3.96	3.82	3.94
4	3	5	10	30	60	5.47	−0.13	5.09	8.21	13.24	8.09	12.89	3.22	3.75	3.80	3.77	3.76
smallest value						3.53	−2.81	−0.36	6.33	10.46	6.21	9.82	3.22	3.71	3.57	3.68	3.55
averages						5.14	−0.24	4.62	7.76	12.50	7.54	12.03	3.75	4.21	3.90	4.21	3.85
largest value						8.33	−0.04	7.90	13.30	16.36	10.86	14.33	4.98	5.20	4.24	5.17	4.18

We consider 25 instances (rows of Table 2.1) that are generated by varying the parameters of our problem G, L, $\check{k}_m$, $\check{h}_m$, K, and H (see columns (a)-(f) in Table 2.1). For example, in the first row of Table 2.1, we have an instance with two groups, the maximum call length is three hours, each group can be called at most five times in a month and at most 20 times over the horizon, and the total time of calls can be at most 10 hours in a month and at most

		Master Problem (aggregate monthly problem; solved on a rolling basis at the beginning of each month)		**Within a Month Problem** (Disaggregation)		Notes
		AMP (stochastic dynamic programming)	**KMP** (multiple Choice Knapsack)	**HWP** (scenario-based lookahead model; solved daily)	**No Disaggregation** (the solution is applied in an aggregate setting)	
Approximation Procedures	**[AMP+HWP]**	☒		☒		• For medium instances.
	[KMP+HWP]		☒	☒		• For large instances.
Bounding Procedures	**[AMP]**	☒			☒	• A lower bound if the monthly decisions and threshold policy are imposed.
	[KMP]		☒		☒	• To measure the error due to using KMP instead of AMP.
	[LB]	This lower bound is estimated through simulation by solving a relaxed deterministic problem for each testing realization and computing the average saving.				• A lower bound on problem P (for any policy). • To measure errors due to all approximations.

Figure 2.8: Summary of our approximation/bounding procedures

40 hours over the horizon. For each parameter setting (each row of Table 2.1), we create 300 annual energy-consumption realizations as the training data. The distribution of energy consumption over the horizon is based on an autoregressive model that we have designed using the actual CAISO data. We use the training data to infer a policy. We then evaluate the policy on a test dataset that consists of a different 300 realizations. Below, we report the performance observed on the test data. Table 2.1 is based on the average savings for the five approximation/bounding procedures listed in Figure 2.8.

Columns (g)-(i) show the performance of our approximations with respect to the "optimal" value of AMP. If the monthly decisions and threshold policy are imposed, the optimal value of AMP is a bound on the optimal value of our problem. Column (g) shows the average percentage loss in saving if the solution of KMP is used for the aggregate problem instead of the solution of AMP. This error changes between 3.53% and 8.33% with an average of 5.14% (see the last three rows of Table 2.1). Columns (h) and (i) measure the difference between the aggregate solution of AMP and the final solution of our approximation method after applying HWP. In other words, columns (h) and (i) show the percentage loss in savings due to disaggregation plus the improvement due to the forward-looking feature of HWP. Compared to AMP, our approximation [AMP+HWP] improves, on average, 0.24% (and between 0.04%

and 2.81%), and [AMP+HWP] loses, on average, 4.62% (and between -0.36% and 7.9%) saving. The largest losses for these approximations are 3.69% and 8.81%, respectively. Note the reason for the negative values in columns (h) and (i) is as follows. Recall HWP allows shorter calls to be skipped in anticipation of longer calls that may arrive in the future, whereas AMP uses all generated calls by DLP (in the aggregate setting). Thus, HWP improves the saving over AMP, and hence, compensates for some portion of the error that is created due to aggregation/disaggregation. Consequently, in some cases, especially when the number of calls is too restrictive, the average saving by [AMP+HWP] is strictly better than that of [AMP], which explains the negative values in columns (h) and (i). Overall, columns (h)-(i) imply the groups-aggregation error is very small.

Columns (j)-(m) measure errors with respect to our LB. Recall LB is a bound on the saving of P obtained by any policy, and consequently, the errors measured based on this bound include the errors due to all approximations as well as the error due to imposing the threshold policy. Column (j) (respectively (k)) could be interpreted as the error due to imposing monthly decisions and the threshold policy (respectively, using KMP instead of AMP and imposing monthly decisions and the threshold policy). Columns (l) and (m) represent the errors of our approximation approaches. Our approximations [AMP+HWP] and [KMP+HWP] produce, on average, 7.54%. and 12.03% errors, respectively. The largest errors we observed among our instances for these approximations are 10.86% and 14.33%, respectively. Note a portion of these errors could be due to the error introduced by the LB because, as we previously stated in §2.3.2, LB relaxes two key constraints and is an ex-post optimization. In short, columns (l) and (m) indicate our approximation procedures provide excellent solutions for this very difficult problem.

Columns (n)-(r) provide the *coefficient of variation* for the percentage saving for the five approximation / bounding procedures. The largest coefficient of variation that we observed was 5.20% (see the last row in column (o)), implying stable performance of these methods across the 300 testing realizations.

Table 2.2: Impact of inputs on the performance of our approximation approaches

changing factor	%error vs AMP			%error vs LB				coefficient of variation (%)				
	KMP	AMP+ HWP	KMP+ HWP	AMP	KMP	AMP+ HWP	KMP+ HWP	LB	AMP	KMP	AMP+ HWP	KMP+ HWP
(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)	(i)	(j)	(k)	(l)	(m)
G	↑↓	↓	↑↓	↓	↓	↓	↓	—	—	—	—	—
L	↑↓	↑↓	↑↓	↓	↓	↓	↓	—	—	—	—	—
$\breve{k}_m,\ \breve{h}_m,\ K,\ H$	↑↓	—	↑↓	↓	↓	↓	↓	—	—	—	—	—
$\breve{h}_m,\ H$	—	—	—	—	—	↓	—	—	—	—	—	—
$K,\ H$	—	—	—	↑	—	↑	—	↓	↓	↓	↓	↓

Next, we discuss how the parameters G, L, k_m, h_m, K, and H impact the performance of our approximations/bounding procedures. Table 2.2 provides a summary of Table 2.1, demonstrating how the errors and coefficient of variations are impacted by an increase in the values of the parameters. The rows in Table 2.2 summarize the corresponding rows in Table 2.1. We make the following observations:

- Increasing the number of groups G, while keeping the other factors constant, decreases the relative errors of our approximations, because of an increase in the magnitude of saving. We did not observe any significant change in the coefficient of variations when G increases.

- Increasing the maximum call length L improves the saving by all five approaches. The reason we observe "↑↓" in columns (b)-(d) is that by increasing L, the saving by AMP improves sharply first and modestly afterwards.

- Increasing $\breve{k}_m,\ \breve{h}_m,\ K$, and H obviously increases the saving by all five approaches.

- Increasing the hours limits $\breve{h}_m$ and H increases the saving under all five approaches, and decreases the relative error of our approximations (columns (e)-(h)).

- Increasing the annual calls and hours limits (K and H) improves the saving under all five approaches, and increases the percentage errors of [AMP] and [AMP+HWP].

In Figure 2.9, we compare the solutions of [AMP+HWP] and [KMP+HWP] on the instance that corresponds to the fifth row of Table 2.1. Figures 2.9(a)-(c) respectively show

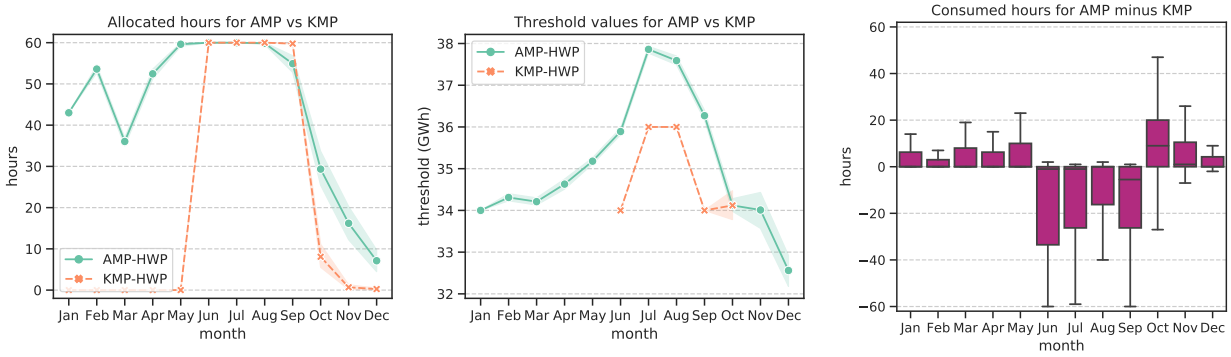

(a) Allocated hours (b) Threshold values (c) Consumed hours

Figure 2.9: Comparing the solutions of [AMP+HWP] and [KMP+HWP] (for the fifth instance in Table 2.1)

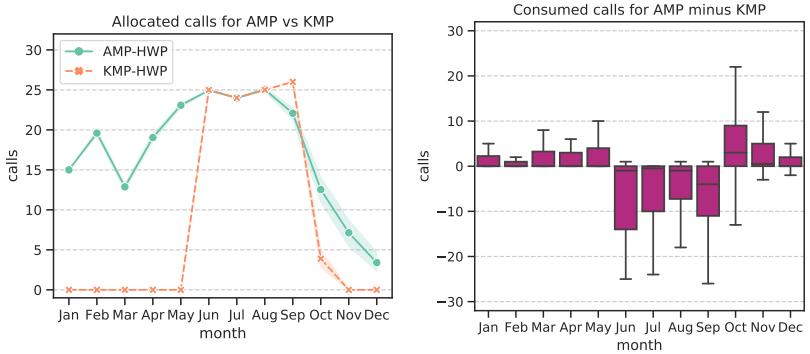

(a) Allocated calls (b) Consumed calls

Figure 2.10: Comparing the number of calls allocated and consumed by [AMP+HWP] and [KMP+HWP].

the total hours allocated to each month (h^*_m), the threshold level for each month ($\mathcal{Z}^*_m$), and the difference in the hours used in each month by these two approaches (the "calls" counterpart of this figure is also presented in Figure 2.10). Note when hours allocated by [KMP+HWP] is zero, its threshold is irrelevant.

As the figures indicate, the [KMP+HWP] approach allocates very small amount of resources in the first half of the year and then increases the allocation to large amounts in the next months. The threshold values that these two approaches use are plotted in figure 2.9b. The threshold values are generally slightly smaller for the [AMP+HWP] approach during the peak consumption months (i.e. June to October). Based on this observation, one would expect that the [KMP+HWP] approach assigns more calls during the peak months. We have calculated the difference between the total number of consumed hours by the [AMP+HWP] approach and the [KMP+HWP] approach and plotted the distribution of these values for each month in figure 2.9c. The plot demonstrates that the average number of hours consumed by these two approaches are similar in all the months. However, we can observe that for some realizations the [AMP+HWP] approach assigns more calls during May and for some realizations the [KMP+HWP] approach assigns more calls during June and September.

As we conjectured in §2.4.3, [AMP+HWP] allocates more resources in the early months (relative to [KMP+HWP]), while it controls the consumption of these resources by employing an appropriate threshold level (see Jan-Apr in Figure 2.9). On the other hand, in the last few month, the allocations, threshold levels, and consumed hours by these two approaches are very similar (see Oct-Dec in Figure 2.9). In summer, and particularly in July and August, which are the hottest months and the main focus of DLCC programs, both approaches have similar solutions. Note [AMP+HWP] controls overflows, and hence has the flexibility to consume some of the resources in the early months. On the other hand, because [KMP+HWP] consumes less early in the year, it compensates in June and September, consuming more than [AMP+HWP] in these two months. Overall, both approaches consume almost all of the resources, although [KMP+HWP] performs this task slightly suboptimally. Recall from Table

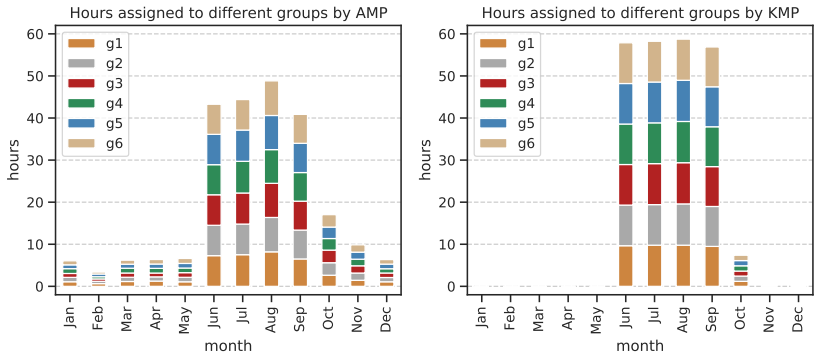

(a) [AMP+HWP] (b) [KMP+HWP]

Figure 2.11: The number of hours assigned to different groups in different months by the two approaches.

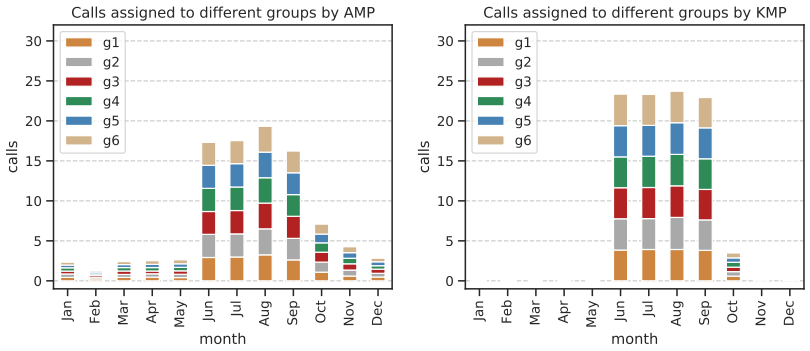

(a) AMP-HWP (b) KMP-HWP

Figure 2.12: The number of calls assigned to different groups in different months.

2.1 that the difference in the amount of saving produced by [AMP+HWP] and [KMP+HWP] is very small.

Figure 2.11 shows the number of hours assigned to each group, in each month, by our disaggregation method HWP. Observe both [AMP+HWP] and [KMP+HWP] maintain a balance allocation of hours across groups throughout the horizon. We observed a similar result for the assignment of calls in Figure 2.12. Thus, our observations and motivation for designing HWP (see §2.5) are confirmed by our numerical analysis, that is, HWP assigns resources to groups in a balanced manner. As we previously stated, a balanced allocation of calls and hours is a fair allocation, which is managerially appealing and desirable for the

DLCC participants.

In summary, we presented the results of our numerical experiments demonstrating the remarkable performance of our approximations. We provided a numerical analysis on how different factors of our problem influence the percentage errors and coefficient of variations. We also presented a comparative analysis of the solution features of our approximation methods.

Next, we present the results of applying our approximation procedures to an industrial instance. We begin by assessing the quality of our solution approach using real data from CAISO. We also report the outcome of implementing our approach at a leading utility firm, which is a member of CAISO, and serves millions of people in California. This work consisted of two phases. In phase 1, we used publicly-available data from CAISO to evaluate our heuristic approach. We measured the potential cost reduction, and performed sensitivity analysis to address the non-monetary concerns of the managers (§2.6.1, §2.6.2, §2.6.3, and §2.6.4). In phase 2, based on the outcomes of an earlier version of our approximation, the firm implemented our approach (§2.6.5). Post implementation, we improved our approximation method based on the recommendations of the review team. Our results in phase 1 are updated and reflect our current approach.

2.6.1 Data collection and parameters estimation

Demand-forecasting models. Utility companies in general, and our partner company especially, have elaborate models for long-term and short-term forecasting of their demand, which is proprietary and can't be disclosed in the dissertation. Consequently, to establish the quality of our solution approach, we developed two simple forecasting models. For long-term planning purposes, we used the historical distribution of the ECPs to simulate the ECPs of the days in each month. We used the model for the aggregate-level planning in AMP and KMP. On the other hand, we developed an auto-regression model to forecast the demand for the next few days. The short-term forecasting model is used in our scenario-based disaggregation

algorithm. We used the CAISO's hourly energy-consumption data from 2014 to 2021 to build and evaluate these models. In the following, we provide more details on the forecast models. In particular, we first briefly explain the long-term forecasting approach that we used for aggregate-level planing and then elaborate on the short-term forecasting model for our scenario-based disaggregation. We acknowledge that there are many alternative sophisticated forecasting models that are available and could be used for both short and long-term estimate of energy consumption.

As explained earlier, we use the historical distribution of ECPs to generate the ECPs of the days for each month. Specifically, for a given day in a certain month, we take random samples from the ECPs of the corresponding month in previous years. By concatenating the generated ECPs for all the days of the year, we then construct a sample for the entire year. This procedure is repeated to generate multiple year-long samples for the aggregate-level planning. Although we should not expect to get precise forecasts for each day, these samples reflect the relative total consumption over different months fairly accurately.

On the other hand, we developed a simple auto-regression model to forecast the demand in the next few days. We used the data from 2014 to 2020 to build a model with the following explanatory variables: (i) month dummies, (ii) weekend dummy (iii) hour-of-day dummies, (iv) one-day and two-day lags of the hourly load consumption. Table 2.3 shows the coefficient of each variable in the regression model. Coefficients for months should be interpreted relative to month 1. As Table 2.3 indicates, hourly electricity consumption rates are significantly higher in June, July, August, and September relative to the baseline (January). Similarly, coefficients for hours should be interpreted relative to hour 1 (00:00 am to 1:00 am). The results suggest the consumption rate is lower at night and has a peak around 8:00 pm. Moreover, the consumption is higher during weekdays. Last, higher consumption today corresponds to higher consumption tomorrow and lower consumption the day after tomorrow. The average and standard deviation of the hourly load are 26,370 MWh and 4,860 MWh, respectively. The in-sample adjusted R-squared of our model is 91.7%, which

demonstrates that the model is an accurate fit to the training data for short-term forecasting. Moreover, step-wise regression determined that all variables are statistically significant at the 1% level. Next, we used the learned model to predict the 2021 hourly consumption rates. Our results indicate that the Mean Absolute Percentage Error (MAPE) on the test data (i.e. out-of-sample) is 4.5%.

Table 2.3: Coefficient of variables in the auto-regressive model

Variable	Coefficient	Variable	Coefficient	Variable	Coefficient
Month (2)	-98.23	Year (2019)	-396.31	Hour (14)	1030.02
Month (3)	-189.35	Year (2020)	-421.22	Hour (15)	1190.44
Month (4)	-133.62	Weekend	-1850.53	Hour (16)	1360.56
Month (5)	121.98	Hour (2)	-266.61	Hour (17)	1549.85
Month (6)	1081.10	Hour (3)	-438.51	Hour (18)	1792.22
Month (7)	1788.71	Hour (4)	-520.89	Hour (19)	1892.79
Month (8)	1819.65	Hour (5)	-477.40	Hour (20)	1864.70
Month (9)	1201.61	Hour (6)	-263.64	Hour (21)	1752.59
Month (10)	309.50	Hour (7)	87.82	Hour (22)	1432.97
Month (11)	-46.69	Hour (8)	368.67	Hour (23)	916.70
Month (12)	86.46	Hour (9)	531.77	Hour (24)	386.34
Year (2015)	-13.06	Hour (10)	629.83	1-Day Lag	0.79
Year (2016)	-69.38	Hour (11)	716.13	2-Day Lag	-0.05
Year (2017)	-49.48	Hour (12)	796.21	-	-
Year (2018)	-190.03	Hour (13)	884.76	-	-

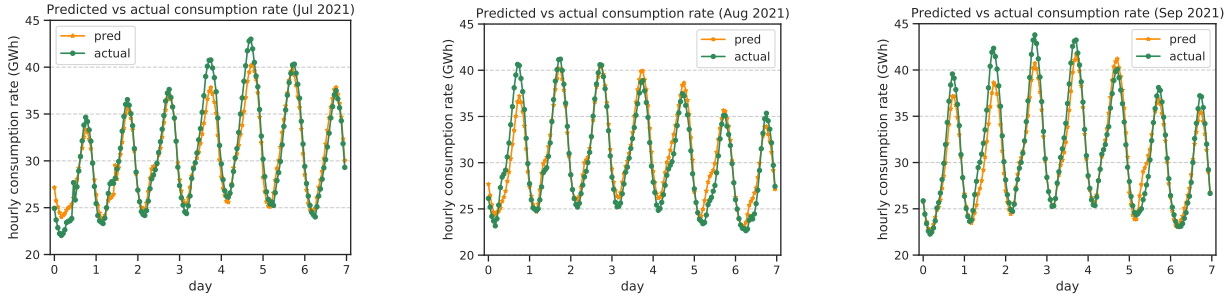

(a) First week of July 2021 (b) First week of August 2021 (c) First week of September 2021

Figure 2.13: Predicted vs actual hourly consumption rates for the first weeks of July, August, and September using the CAISO 2021 data.

In Figure 2.13, we plot the predicted versus actual consumption rates for the first weeks of July, August, and September 2021 using the CAISO data. Although our model is generally fairly accurate, it sometimes under predicts or over predicts a group of samples that are associated with large consumption rates. This might be due to the lack of sufficient number

of observations from these cases or related to the inherent difficulty of predicting high-consumption events. By taking a closer look at the prediction time series data, one can observe that our model sometimes misses daily trends when there are sudden positive or negative jumps in the consumption rate from one day to the next day. Since one-day lag is the most recent piece of information that is being used by our model for making predictions for any given day, promptly catching up with these jumps is generally a challenging task. However, the results indicate that the model always corrects its prediction at most one day after such drastic changes in the consumption trend occur.

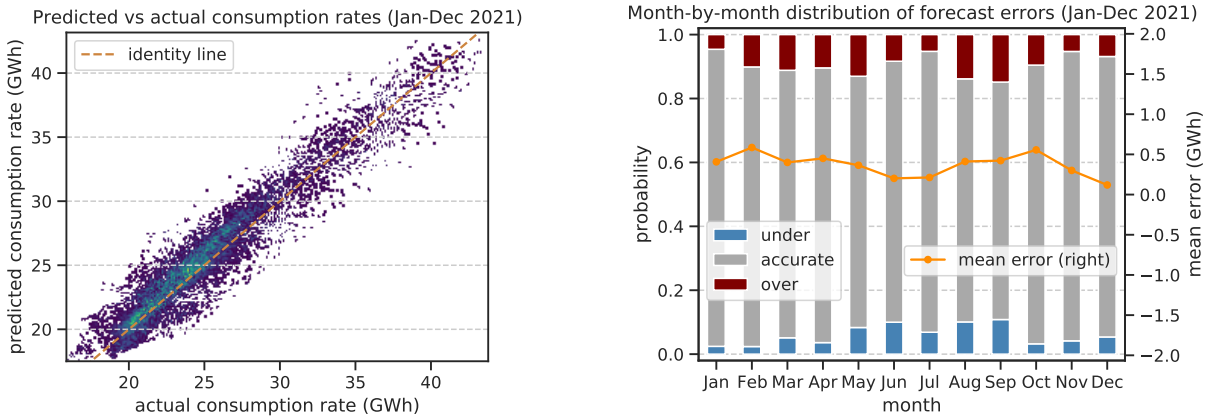

(a) Predicted vs actual scatter plot

(b) Month-by-month breakdown of forecast errors

Figure 2.14: The first figure presents the scatter plot of predicted vs actual hourly consumption rates using the CAISO 2021 data (lighter colors correspond to higher density of observations). The second figure plots the month-by-month probabilities of large under-prediction and over-prediction by the model (left axis) as well as mean forecast error (right axis) using the same data. Accurate prediction is defined as having an absolute forecast error of at most 2 GWh.

In Figure 2.14a, we present the scatter plot of actual versus predicted consumption rates for all the test samples (January to December 2021). Furthermore, the month-by-month probabilities of large under-prediction and over-prediction of hourly consumption rates by the model, as measured by having 2 GWh or more forecast error on the test data, are presented in Figure 2.14b. Notice that the threshold level has been determined in a way that reflects concerns of practitioners regarding the forecasting accuracy of the model. In addition,

we show the mean monthly forecast errors on the same plot. As the figures indicate, the probability of accurate forecast on unseen data, even in the worst case (i.e. September), is still above 70%, which is acceptable from a practical perspective. Even though the model may systematically over-predict or under-predict the consumption rates on unseen data due to a variety of factors, we observe that the mean forecast errors are very close for different months.

Cost-function estimation. To estimate the cost function, we fitted a quadratic curve to the marginal generation cost [NSP12] as a function of total capacity with a zero intercept. This fit resulted in cost function $f(z) = 0.0229z^3$, where z denotes the load in GWh (note the total cost function is of degree 3 because the marginal cost is quadratic). The R-square of the fit was 0.8.

2.6.2 Comparison with prior practice

We compare the performance of our approach with the firm's prior practice on CAISO data. This step was crucial in demonstrating the financial benefits of our model and convincing the firm to adopt our approach. The DLCCs that firm offers allow up to one call per day, 10 calls per month, and 100 calls per year for each group. The total hours are limited to 20 hours per month and 180 hours per year. Most of the firm's service areas have 20 customer groups. Because the firm expects to expand its existing DLCCs or offer them in new service areas in the future, it asked us to consider 10 to 25 groups in our experiments.

The firm's prior practice. The firm's prior practice employed a single fixed threshold over the entire year. For confidentiality reasons, we are unable to disclose the threshold. We are only allowed to state that it would be equivalent to 36 GWh if the practice were to apply throughout CAISO. At the beginning of each day, if the forecast ECP indicated the peak load would exceed the threshold, the firm would generate and assign calls to groups to reduce the consumption to the fixed threshold. Although the contracts stated that a single call could last for up to six hours, an internal tradition sought to achieve the target level by using calls

that were *exactly* two hours long. If two-hour calls were not sufficient, the firm used calls that were three hours long. The main reason the firm preferred to use two-hour or three-hour calls was to save the longer calls for reducing load during emergency events (see §2.6.4).

The firm used a rotational policy, as that in DA, to call groups. Groups were called in the order of 1 to G recursively. At the beginning of each day if the number of two-hour calls for reducing peak load to 36 GWh is less than or equal to the number of groups and the number of groups that have at least one remaining call and 2 remaining hours is at least equal to the number of calls to be made, then the calls are assigned to groups. The available number of calls and hours for each group in the current month and in the year are updated. On the other hand, if the number of 2-hour calls exceeds the number of groups, then the algorithm checks whether the number of 3-hour calls is less than or equal to the number of groups and the number of groups that have at least one remaining call and 3 remaining hours is at least equal to the number of 3-hour calls. If this is the case, then the 3-hour calls are assigned to groups.

Comparison with the firm's prior practice. To make a fair comparison, we run our model with $L = 2$. Table 2.4 shows the additional percentage reduction in cost when we use our model instead of the firm's prior approach. For example, for the instance with 10 groups, our suggested approach reduces the total cost by 1.79% over the utility's prior practice. Despite excluding the possibility of using three-hour calls, our solution approach could have reduced the total cost by 3.11%, on average, and as much as 4.70%. We repeated this experiment with $L = 3$ and obtained similar results. Intuitively, when $L = 2$, the constraints on calls and hours to each group still allow our model to use shorter (compared to $L = 3$) but more calls to reduce the peak load to optimal thresholds in each month. Note our solution only uses calls that are at most two hours long, but the prior practice had the option to use three-hour calls. The capital-intensive nature of the utility industry makes these savings highly valuable. Additionally, if we account for other financial benefits of reducing the peak-load consumption, such as reducing the number of blackouts and abating greenhouse gas emissions and pollution

from electricity generators, the total benefit is much larger.

Table 2.4: Our approximation vs the firm's prior practice ($L = 2, \breve{k}_m = 10, \breve{h}_m = 20, \forall m, K = 100, H = 180$)

number of groups	10	15	20	25	average
cost savings	1.79%	2.47%	3.49%	4.70%	3.11%

We also investigated the effect of using a single threshold throughout the horizon versus dynamically updating the threshold each month. We modified our algorithm to select one threshold for the entire year, at the beginning of the year, and employ that same threshold throughout the year. In our industrial instances with 10 and 20 groups, the expected saving decreased by 6.15% and 5.42%, respectively. Thus, we conclude a large portion of our saving is due to appropriately determining a threshold value as well as the effectiveness of our disaggregation procedure HWP, whereas a relatively smaller portion of our saving is due to customizing the threshold for each month. The value of customizing the threshold for each month increases if the number of groups decreases, monthly constraints become tighter, and the consumption significantly changes with-in a month.

2.6.3 Risk analysis

In addition to expected costs, managers are also concerned about the risk of adverse events. Measures of risk that our partner utility company monitors are (a) likelihood of power consumption exceeding some critical level, (b) the number of days in which the power consumption exceeds a given critical level, and (c) the expected total power generated above a critical level. Minimizing costs may not result in minimizing these measures of risk due to uncertainty in electricity demand. Recall we allocate resources to each month and set a threshold level for the month in the aggregate-level planning. However, the threshold level that minimizes the expected cost may not be the same as the threshold level that minimizes risk. Intuitively, a higher threshold level may decrease the likelihood of having higher peak loads, because fewer resources will be consumed for cutting lower consumption levels.

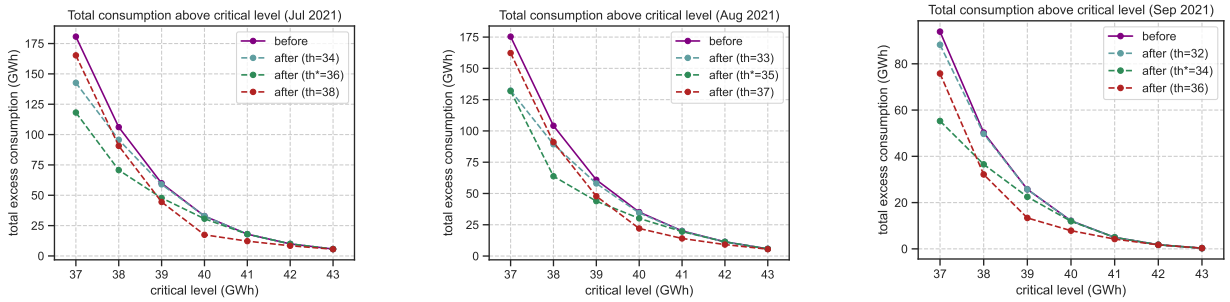

(a) July 2021 (b) August 2021 (c) September 2021

Figure 2.15: Total energy consumption above different critical load levels during peak months for different choices of threshold.

Figure 2.15 illustrates our risk analysis. As we stated earlier, we consider the total energy generated when the system load is above a critical level as a measure of the risk of network failure. To quantify the risk of our approach, we generated a set of sample ECPs for 2021 using historical data from 2014 to 2020, and applied our approach to each sample to allocate calls and hours over the entire year. Then, we computed the total energy consumption above different critical levels for each sample. Figure 2.15 shows the average over the samples for the peak months. The horizontal axis indicates the critical load levels and the vertical axis is the total power generated above these critical levels during the month. We compare the risk before and after applying our approach, for different critical levels (asterisks denote the cost-minimizing thresholds found by our approach).

As expected, the total consumption above all critical levels decreases after applying our approach. We also observed that manually increasing the threshold level above the cost-minimizing level increases the odds of shaving off higher levels of consumption, which decreases the total consumption above higher critical levels. In other words, the risk of network failure due to high system demands decreases when we slightly increase the cost-minimizing threshold level. Indeed, as a drawback, choosing a higher threshold level increases the cost of generating electricity. The experimental results demonstrate that tightening (decreasing) the threshold doesn't offer any advantage as the total consumption measure is increased for high

cutoff values.

Our approximation (KMP+HWP) extends to allow the managers to obtain a right balance between the expected cost and risk, by leveraging the monthly threshold levels. We extend KMP as follows. In month m, let a_m denote a given acceptable expected total consumption above some critical value $\mathcal{Z}_m$. When we construct our KMP instance, we determine the monthly costs $\Phi_m(k, h)$ as follows. Given k and h, we select a cost-minimizing threshold from all potential threshold levels that have at most a_m expected consumption above the critical value. Note if the utility requires a too small a_m, no threshold may achieve this goal. In this case, the utility may consider increasing the participants and/or resources (calls and hours). The monthly threshold that we use as an input for HWP also needs to be selected from all potential thresholds that have at most a_m hours of violation. Other components of our methodology readily extend to incorporating the risk of exceeding thresholds. Furthermore, our approximation similarly extends if the utility wishes to use other risk measures such as the probability of exceeding a critical value.

2.6.4 Managerial considerations

Based on the performance of our model, the utility firm's director of the DR program expressed interest in adopting the model. But before proceeding, the director asked us to investigate a few additional issues. DLCCs are a new type of DR program, and their growth depends on customers' willingness to sign up and authorize utility firms to interrupt their consumption. Although financial incentives play an important role in customers' adoption of DLCCs, other factors potentially influence customer participation. The experiments that we report next aimed to address the critical concerns of the managers that mostly focused on enhancing participating customers' experience and encouraging more participation in the DLCC programs. In the remainder, we list these managerial concerns and discuss how we addressed them.

Fewer but longer calls. Some of the managers were concerned about the number of calls,

and suggested that the firm uses fewer but longer calls. They believed the number of calls (monthly and annual) could be annoying to existing customers and result in a negative word of mouth about the DLCCs. They were curious to see the effect of reducing the number of calls by a factor of 10% to 30%. To address this scenario, we ran new experiments with our model using $L = 6$, $breve{h}_m = 20, \forall m$, and $H = 180$. We also considered 10 and 20 groups. For monthly and annual calls, we used $\breve{k}_m = \lfloor 10(1-\varepsilon) \rfloor, \forall m$, and $K = \lfloor 100(1-\varepsilon) \rfloor$, with ε denoting the reduction factor and $\lfloor \cdot \rfloor$ denoting the floor function. We tried 5%, 10%, 15%, 20%, 25%, and 30% for ε. Interestingly, we found reducing the monthly and annual number of calls from their base values (10 and 100) does not decrease the savings. With $\breve{h}_m = 20, \forall m$, and $H = 180$, allowing for calls of up to six hours (vs. 2 or 3 hours) compensates for the reduction in the number of calls, even when the reduction is 30%. Therefore, if the firm ultimately decides to use fewer but longer calls, doing so without sacrificing the savings is possible. In other words, given the aforementioned values of L, H_m, H, and G, the number of calls (K_m and K) are not binding, even if their values are reduced by a factor of 30%.

Imposing minimum call duration. Some managers were in favor of not disturbing customers for one-hour calls, and suggested using a lower limit of two hours on the duration of calls. We proposed two ways of incorporating this lower limit in our approximation. First, if the length of a call in Algorithm 1 is less than two hours, the call continues until it reaches the minimum duration of two hours. In this approach, the load could drop below the optimal threshold in some hours. Second, if the duration of a call does not meet the minimum duration, the call is not generated in the first place. In this approach, the load may exceed the optimal threshold in some hours. After proposing these approaches, the managers asked us to use the first approach. We ran our experiments with $L = 6$, $\breve{k}_m = 10, \breve{h}_m = 20, \forall m$, $K = 100, H = 180$, and $G = 20$, adding a constraint that calls should be at least two hours. The results showed a small change of 0.02% in cost. Thus, we concluded that, given the parameters of the firm's DLCCs, imposing a lower limit of two hours while using $L = 6$ does not have a significant effect.

Reserving calls for emergencies. Some managers strongly felt that they should not use all the calls and hours. Rather, they should reserve some of the calls and hours for managing the peak load during unpredictable and irregular events, such as a sudden shock to the system. The managers' recommendation was to keep the same number of calls and hours in the contract, but internally reserve 10%, 15%, or even 20% of the calls and hours. To examine the effect of this strategy on cost, we repeated our experiments with $L = 6$, but this time we used $\breve{k}_m = \lfloor 10(1-\varepsilon) \rfloor, \breve{h}_m = \lfloor 20(1-\varepsilon) \rfloor, \forall m$, $K = \lfloor 100(1-\varepsilon) \rfloor$, and $H = \lfloor 180(1-\varepsilon) \rfloor$, with ε equal to 5%, 10%, 15%, and 20%. This way, 5% to 20% of calls and hours were reserved for rare events. Our numerical analysis indicated that if only less than 10% of calls and hours are reserved, then cost increases by between 0.11% and 0.71%. On the other hand, if the firm wants to be more conservative and reserve 15% to 20% of calls and hours, then cost will increase by more than 1% (between 1.31% and 1.47%) when there are 20 and 25 groups. The reason cost increases for $G = 20$ is the same for 0.15 and 0.2 is as follows: when $\varepsilon = 0.15$, there are several groups with non-zero slacks for calls and hours in the optimal solution. When ε increases to 0.2, the algorithm assigns additional calls and hours to the groups with slacks and it can achieve the same reduce-to thresholds as when $\varepsilon = 0.15$. Thus the increase in cost is the same for 0.15 and 0.2. However, at $\varepsilon = 0.2$, the annual hour constraint becomes binding for all the groups, so the curve will start to increase if ε increases beyond 0.2. Calls and hours that are not used in a given month do not carry over to future months, so the firm forfeits the value of these tokens if it ends up not using them. Therefore, the firm should select the level of reservation carefully and not be overly conservative.

Fair allocation of calls to groups. Some managers were interested in seeing the distributions of total calls and hours made to groups over a year, for fairness considerations. Figures 2.16 and 2.17 show the histograms of the annual calls and hours assigned to 20 groups under our model and the firm's prior practice. Under the firm's prior practice, the average number of annual calls is 27, the standard deviation is 1.91 calls, and the majority of groups receive 28 calls in a year. The difference between the maximum and minimum number of calls is

14. Under our model, the average number of annual calls is 36.71, the standard deviation is 1.12 calls, and the majority of groups receive 37 calls in a year. The difference between the maximum and minimum number of calls is 7. Naturally, our algorithm makes more calls because it uses the calls and hours to reduce the load in each month to the threshold that would minimize the total cost. However, the distribution of annual calls in our model has a lower coefficient of variation.

Under the firm's prior practice, over a year, the average and standard deviation of annual hours are 54 and 3.81, respectively, and the difference between the maximum and minimum hours 28. The distribution of annual hours in our model has an average and standard deviation of 161.94 and 9.73 hours, respectively, and the difference between the maximum and minimum hours is 60. Thus, the distribution of annual hours has a higher range, but lower coefficient of variation compared to the prior practice. As a result, the coefficient of variation is lower than the firm's practice. Note our model uses more hours because it uses calls longer than two and three hours (up to 6 hours) if necessary- and targets thresholds that may be below the firm's previously used threshold.

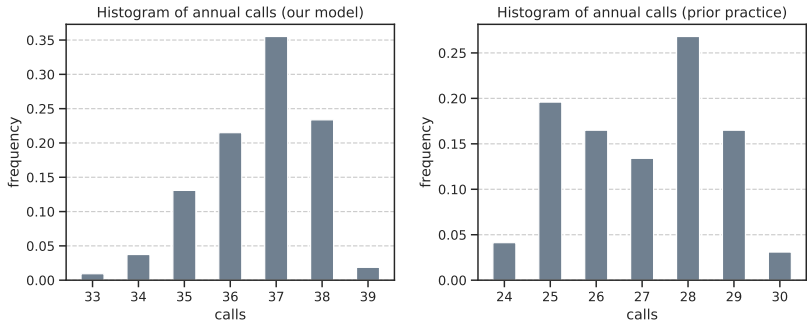

(a) Our model (b) Prior practice

Figure 2.16: Histograms of annual calls ($G = 20, L = 6, k_m = 10, h_m = 20, \forall m, K = 100, H = 180$).

We also compared the number of days between consecutive calls to the same group to see how frequently our model would call groups. Figure 2.18 shows the histograms of the number of days between consecutive calls to the same group. Under the firm's practice, the average

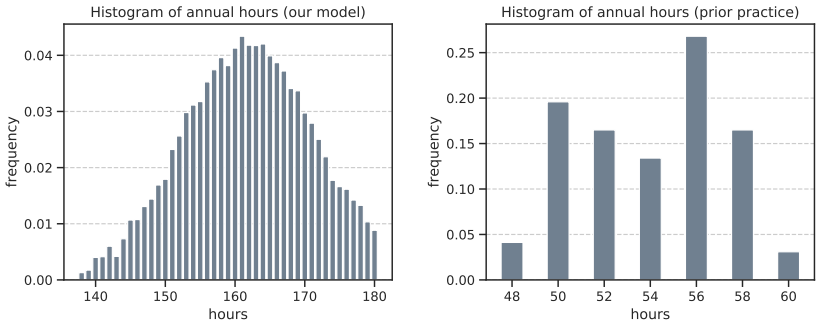

(a) Our model (b) Prior practice

Figure 2.17: Histograms of annual hours (same parameters as Figure 2.16).

and standard deviation of days between consecutive calls are 3.66 and 2.57, and the majority of the groups are called every two days. On the other hand, our model calls groups every 9.42 days, on average, with a standard deviation of 3.41 days. Also, most groups are called every nine days. Thus, our model calls customers less frequently. Recall an advantage of our approach is to assign calls to groups while minimizing the imbalance among groups, which results in fair allocation of these calls. Our above numerical studies confirms this aspect of our approach.

Figures 2.16, 2.17, and 2.18 together assured the managers that the load-reduction policies of our model do not lead to excessive variability in calls and hours across the groups and that although our model uses longer hours, it interrupts customers' energy consumption less frequently.

Load-shifting behavior. DLCCs contractually prohibit customers from deliberately shifting the load from peak to off-peak periods. The punitive clauses in contracts dissuade most customers from load shifting, and the technology that is used in remote control devices enables utilities to detect violations. Nevertheless, the managers believed some customers would still shift the peak load. They conjectured that 10% of customers shift 10% to 20% of the peak load to off-peak periods. The managers were interested in measuring the impact of this violation on cost. We modeled this behavior in our experiments and observed that load shifting did not result in off-peak periods becoming new peak periods. Because higher consumption during

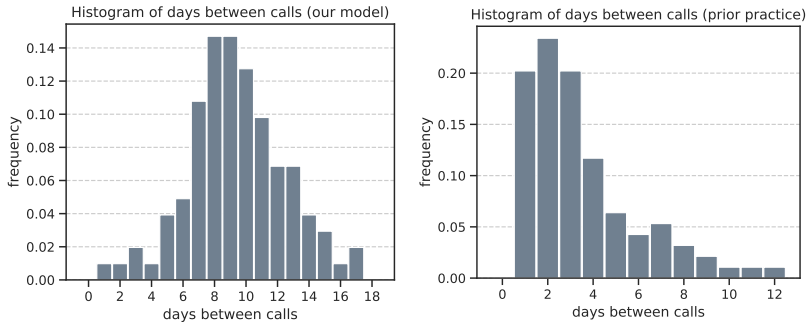

(a) Our model (b) Prior practice

Figure 2.18: Histograms of days between calls (same parameters as Figure 2.16).

off-peak periods generates revenue for the utility firm, the firm does not need to be concerned as long as load shifting does not create a new peak period. If load shifting results in the creation of new peak periods, then the firm could penalize customers who violate the contract. Note if load shifting to off-peak periods is allowed, our approximation approach could be used as long as it does not generate a new peak.

Shorter call lengths on hot days. Some managers suggested that calls be limited to a maximum of two or three hours during hot summer days, to avoid the potential negative impacts of having to reduce load at extreme heat. Because reducing the maximum duration of daily calls may result in a larger number of calls or longer hours over a year, we were asked to check the effect of using a variable maximum call duration on total cost as well as the average number of calls and hours over a year. To address this concern, we considered three scenarios: when $L = 2$ in summer (July to September) and $L = 6$ otherwise, when $L = 3$ in summer and $L = 6$ otherwise, and when $L = 6$ throughout the year (the benchmark). We found reducing L in the summer months to two or three did not increase the cost. We observed that varying L has a minimal effect on the average annual hours, but it increases the number of calls received over a year. Restricting the duration of summer calls to three hours maximum results in a 7.82% (for 10 groups) to 14.32% (for 20 groups) increase in the average number of annual calls. A further restriction of calls to two hours in the summer increases the average number of annual calls by 26.52% (for 10 groups) to 33.24% (for 20 groups). Thus,

placing a tight restriction on the duration of calls can actually backfire; although limiting the maximum call length alleviates the negative impact of long daily interruptions, customers may end up receiving more calls, thereby creating a different negative impact. Therefore, the firm should pursue a moderate reduction in the maximum call duration in summer months.

Setting contract terms. Fine tuning the contract terms is an important managerial concern. To capture the effect of contract terms, we repeated our industrial instance with 10 groups for the cases of 10% fewer/more calls/hours. In Table 2.5, each cell shows the change in total saving relative to the base case. For example, if we have 10% fewer calls and hours, the total saving reduces by 10.37%. Using this table, we observe a change in the number of hours is more consequential than a similar change in the number of calls. This table is a valuable decision support tool for effectively designing DLCCs.

Table 2.5: Marginal change in saving ($G = 10$)

	-10% hours	base values	+10% hours
-10% calls	-10.37%	-0.16%	5.41%
base values	-8.26%	0.00%	9.52%
+10% calls	-8.21%	0.00%	9.77%

Table 2.5 can also be used to understand the consequences if the firm wishes to keep the same contract terms but internally reserves some of the calls and hours for managing the peak load during unpredictable and irregular events, such as a sudden shock to the system; for example, reserving 10% of calls and hours decreases the expected saving by 10.37%. Last, our numerical results indicates reducing L from six to four hours does not have an impact on saving. Fixing all calls to be three hours long decreases the saving by 5.22%.

The impact of demand fluctuation. Higher volatility in demand due to climate change will increase the value of DLCCs. Higher volatility implies higher peak consumption levels. CAISO reported that on September 6^{th} 2022 the energy consumption set a new record at 52 Gigawatts. The utilities sent text messages to all users requesting them to reduce power. This proved to be highly effective. To illustrate the impact of higher volatility, we conducted

additional experiments and presented the results below.

We use shifted historical distribution of the ECPs to simulate our training and testing realizations. The simulations are based on our short term forecasting model. To see how a higher/lower fluctuation impacts our solution, we use a small example with $G = 4,\ L = 3,\ \breve{k}_m = 5,\ \breve{h}_m = 10,\ K = 20$, and $H = 40$, and repeat our experiment with different noise levels. We let the standard deviation of the random noise change from 40% to 160%, in the increments of 20% (100% fluctuations refer to the base value). Figure 2.19 shows the result of this experiment. Each data point is the average of five randomly generated instances, where each instance has 300 training and 300 testing realizations. In this figure, we let the expected saving of [KMP+HWP] with a noise level of 100% be the base value. We normalize all expected saving values relative to the base value (100% fluctuations); for example, when fluctuation is 160% (i.e., 60% higher than the base case), the expected saving of [KMP+HWP] is 4.51% higher than the base value. We observe that the saving increases when the fluctuation increases. Hence, DLCCs are more advantageous when the demand patterns are more volatile. Furthermore, we observe the expected saving of LB increases faster than that of KMP and [KMP+HWP]. This result is simply because in our lower bound we consider each realization separately and assume we know with certainty the consumption throughout the year. Thus, obviously, increasing fluctuations significantly benefits the lower bound. Simply put, as the noise level increases, the lower bound becomes looser. Therefore, the increased difference between LB and [KMP+HWP] does not necessarily mean our approximation error increases when the fluctuations increase.

2.6.5 Real-world implementation

The approval to implement our approximation method was based on the results of our numerical experiments as well as the analysis of the CAISO data, which clearly indicated superiority of our approach to the prior ad-hoc method. Later in the implementation, we used the firm's data, their adapted demand models, and the firm's historical demand and

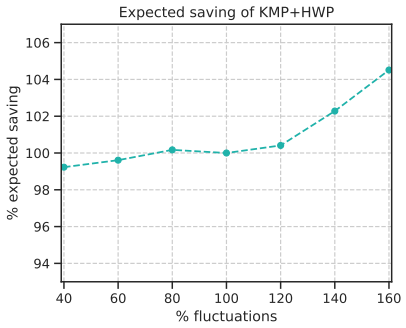

Figure 2.19: Impact of higher/lower fluctuations

desired threshold levels. We discuss some of the issues and implementation complications below.

As explained earlier, the firm has its own short-term and long-term demand-forecasting models, which are much more sophisticated than the simple models that we used in our experiments. In addition to employing various data sources and incorporating the domain knowledge of their forecast team, a key difference between their forecast models and ours is that the firm's managers are more concerned about the capability of the model to accurately predict the peak-load consumption. The modular design of our algorithm allowed us to conveniently replace our basic model with the firm's model in the actual implementation phase.

The second complexity was the time of notice of calls to customers and its implications on demand-forecasting accuracy and customer compliance. Although our model could use and benefit from the most recent estimate of the demand, generated in the early morning, and schedule the calls based on that, some managers did not authorize the "early morning" notifications of the calls, because they had a strong preference to give more time to subscribers to prepare for the scheduled events. Hence, giving "one-day notice" became the preferred mode of operation.

Additionally, the managers needed to approve the scheduled events prior to dispatching them to the customers. In general, some managers were reluctant to approve long calls (five-

and six-hour calls) to customer groups, particularly when the peak-load demand was not much higher than the desired threshold consumption level. We were informed that about 11 percent of the long calls, five and six hours, were not scheduled as recommended and were replaced with four-hour calls. To incorporate the management overriding policy, the schedulers used an ad-hoc approach to schedule additional calls to compensate for these changes, which resulted in more calls being scheduled, as expected.

Customer compliance was another complication in our implementation. Although the data showed that at most 4% of the customers did not comply with the scheduled calls, the variability in customer compliance was a concern. Our model does not consider the variability that compliance introduces. In section §2.7, we consider stochastic customer compliance with scheduled calls.

Despite some discrepancies between the optimal events scheduled by our model and the actual events used after management approval, our model helped the firm reduce costs significantly. We were informed the additional reduction in cost after implementation of our model was approximately 4%. As we mentioned earlier, the total saving increases even further if we add the medium- to long-term benefits of curtailing energy consumption during peak periods, such as the positive impact on the environment and a higher reliability of the grid.

2.7 Model Extension: Stochastic Customer Compliance

We extend our analyses to incorporate stochastic customer compliance behavior and discuss its implications for our solution approach. First of all, notice that participating customers in DLCC programs are contractually obliged to comply with the ”calls” assigned by the utility companies. These contracts usually contain a penalty clause for not fully adhering to the calls. Moreover, utility companies usually send couple of notices to increase the likelihood of customer compliance. The historical data provided to us by our partner utility company indicate 97.4% of the time the customers fully comply with the assigned calls.

Consider scheduling a call to a group with start and end times t' and t'', respectively. Recall, assuming full compliance, the total reduction in consumption in period $t \in \{t', \ldots, t''\}$ due to this call is 1, recalling we have normalized the reduction per group to 1. To model compliance uncertainty, we assume the compliance in period t, denoted by χ_t, has an arbitrary distribution $\mathcal{F}_\chi(\cdot)$ with mean μ, for example, uniform distribution $\chi_t \sim \mathcal{U}[0.8, 1]$. We assume $\mathcal{F}_\chi(\cdot)$ is identical across different periods, groups, and calls. Under these assumptions, a call from t' to t'', in expectation, reduces the total consumption by μ in each period $t \in \{t', \ldots, t''\}$.

Although we assume χ_t has a uniform distribution, one may also use a truncated normal distribution by, first, generating a normally distributed number with mean μ and standard deviation ς, and, then, truncating it to be between 0 and 1. On the other hand, the Beta distribution can also be used as a more natural alternative than the truncated normal distribution for modeling the compliance probability. In this case, we set $\alpha = \left(\frac{1-\mu}{\sigma^2} - \frac{1}{\mu}\right)\mu^2$ and $\beta = \alpha\left(\frac{1}{\mu} - 1\right)$ to get a Beta distribution with mean μ and standard deviation σ. We have repeated our numerical experiments with several alternative distributions and our observations were consistent with those for the uniform distribution.

Let $f(.)$ denote the hourly cost function (see Appendix 2.9), and assume this function is identical across different periods. Consider an arbitrary period t with a known initial consumption $r \in \{0, 1, \ldots, Z\}$. Let $k \in \{0, 1, \ldots, G\}$ calls be active in this period; that is, k groups are on call in period t. The i^{th} call ($i \in \{1, \ldots, k\}$) reduces the total consumption by χ_i. Thus, the expected cost in period t is

$$\varpi_{r,k} := \mathbf{E}_{\chi_1,\ldots,\chi_k}\left[f\left(r - \sum_{i=1}^{k} \chi_i\right)\right].$$

We use simulation to estimate $\varpi_{r,k}$ for all $r \in \{0, 1, \ldots, Z\}$ and $k \in \{0, 1, \ldots, G\}$. We simulate 300 scenarios for $\sum_{i=1}^{k} \chi_i$ to estimate $\varpi_{r,k}$ in our numerical experiments. These simulations are performed in a reasonable amount of time for the chosen distributions of $\mathcal{F}_\chi(\cdot)$.

In the following, we present a lower bound and argue a greedy algorithm can be used for

calculating its value. Moreover, we explain estimating the monthly costs $\Phi_m(k,h)$, which is the key component of KMP that is affected by allowing stochasticity in the compliance rate. We also discuss extending Algorithm 1 to calculate daily calls $\{l_{d,s,i}\}_i$, which affects our loss function $\Psi_{m,d}(\cdot)$ in HWP. First, we discuss how to extend the lower bound algorithm. Under stochastic compliance, we compute our lower bound as:

$$
\begin{aligned}
LB \;&:=\; \mathbf{E}_{\{\tilde{r}_{mdt}\}_{m,d,t}}\left\{ \min_{\substack{\sum_{m,d,t}\mathcal{R}_{mdt}\le G\min\{H,LK\},\\ \sum_{d,t}\mathcal{R}_{mdt}\le G\min\{\breve{h}_m,L\breve{k}_m\},\ \forall m,\\ \sum_{t}\mathcal{R}_{mdt}\le GL,\ \forall m,\ d,\\ \mathcal{R}_{mdt}\in\{0,1,\ldots,G\},\ \forall m,\ d,\ t,}} \left\{ \sum_{m,d,t} \mathbf{E}_{\chi_1,\ldots,\chi_{\mathcal{R}_{mdt}}}\left[f\left(\tilde{r}_{mdt} - \sum_{i=1}^{\mathcal{R}_{mdt}} \chi_i \right) \right] \right\} \right\} \\
&=\; \mathbf{E}_{\{\tilde{r}_{mdt}\}_{m,d,t}}\left\{ \min_{\substack{\sum_{m,d,t}\mathcal{R}_{mdt}\le G\min\{H,LK\},\\ \sum_{d,t}\mathcal{R}_{mdt}\le G\min\{\breve{h}_m,L\breve{k}_m\},\ \forall m,\\ \sum_{t}\mathcal{R}_{mdt}\le GL,\ \forall m,\ d,\\ \mathcal{R}_{mdt}\in\{0,1,\ldots,G\},\ \forall m,\ d,\ t,}} \left\{ \sum_{m,d,t} \varpi_{\tilde{r}_{mdt},\mathcal{R}_{mdt}} \right\} \right\}
\end{aligned}
$$

where $\tilde{r}_{mdt}$ is the initial consumption in period t of day d of month m, in a given testing realization, and $\mathcal{R}_{mdt}$ denotes the number of active calls in this period. For any $\mathcal{R} \in \{0,1,\ldots,G-1\}$, the marginal change $\varpi_{\tilde{r},\mathcal{R}} - \varpi_{\tilde{r},\mathcal{R}+1}$ is non-increasing in $\mathcal{R}$. Therefore, the minimization problem in our lower bound is solved using a greedy algorithm as follows. At an iteration, assume so far we have assigned $\mathcal{R}^{\circ}_{mdt}$ to period t of day d in month m. We compute the marginal change $\varpi_{\tilde{r},\mathcal{R}^{\circ}} - \varpi_{\tilde{r},\mathcal{R}^{\circ}+1}$ for all available periods (i.e., the periods to which assigning an extra call is feasible) throughout the horizon, and assign one call to the period with the largest marginal change. We repeat until no available period is left.

The approximation scheme of [KMP+HWP] can also be conveniently extended to the stochastic compliance setting. The only impacted element of KMP is $\Phi_m(k,h)$, which denotes the expected cost in month m if k calls and h hours are available. To estimate $\Phi_m(k,h)$, for each potential threshold $\mathcal{Z}$, we need to compute the expected cost for each training realization

of month m. For a fixed threshold $\mathcal{Z}$ and a given day in a training realization, we estimate the expected cost as follows. When we schedule calls, we approximately assume each active call deterministically reduces the consumption by μ. We schedule as many calls as needed to reduce the consumption below threshold $\mathcal{Z}$ (in expectation). For a fixed threshold $\mathcal{Z}$, in a given period t with initial consumption r, because we assume each active call deterministically reduces the consumption by μ, we can compute the number of calls that must be active in period t to reduce the consumption r below $\mathcal{Z}$. Therefore, for a given ECP and threshold $\mathcal{Z}$, the total hours is fixed and computed as $\hat{h}_{md} := \sum_{t=1}^{T} \lceil \frac{[r_{mdt}-\mathcal{Z}]^+}{\mu} \rceil$. To compute the number of calls, we use Algorithm 1 except that in line 5, we update the load as $\bar{r}_t := \bar{r}_t - \mu$. After we schedule calls, we compute the expected cost by incorporating the stochasticity around μ; if the initial consumption in period t is r and we have scheduled $\mathcal{R}$ calls to be active in this period, the expected cost in period t is $\varpi_{r,\mathcal{R}}$. We repeat this procedure for each potential $\mathcal{Z}$ and select the one that has the smallest expected cost for the month. After estimating $\Phi_m(k,h)$, for all m, k, and h, we solve KMP.

By solving KMP, we determine the total calls k_m^* and hours h_m^* to be used in month m as well as the threshold $\mathcal{Z}_m^*$. To deploy these calls and hours in month m, we solve HWP on a daily basis. The only impacted elements of HWP are $\{l_{d,s,i}\}_i$ and $\Psi_{m,d}(\cdot)$. Similar to our above explanation, we use modified Algorithm 1 to determine $\{l_{d,s,i}\}_i$. Note the input for estimating $\Psi_{m,d}(\cdot)$ is $\{l_{d,s,i}\}_i$; that is, $\Psi_{m,d}(\cdot)$ is impacted through $\{l_{d,s,i}\}_i$. In short, HWP extends to the case of uncertain compliance simply by applying the modified Algorithm 1 to determine $\{l_{d,s,i}\}_i$.

Numerical experiments. We solved our 25 simulated instances (introduced in §2.6) using a uniformly distributed compliance between 65% and 85%. The average error of KMP+HWP increased from 12.03% (for the case of deterministic full compliance, recalling from the first paragraph of §2.6) to 12.21%. Thus, we conclude our generalized approximation for the case of uncertain compliance performs remarkably well.

We performed additional experiments using a small example with $G = 4,\ L = 3,\ \breve{k}_m =$

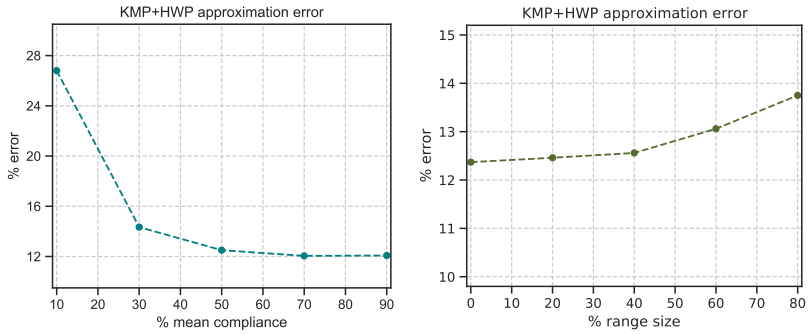

(a) KMP+HWP error vs mean compliance

(b) KMP+HWP error vs range size

Figure 2.20: Performance of our approximation with uncertain (uniformly distributed) compliance. The approximation error is measured as % error in comparison to the lower bound. In figure (a), we set the range size to 20%. Similarly, we set mean compliance to 50% in figure (b).

5, $\breve{h}_m = 10$, $K = 20$, and $H = 40$, to investigate the impact of the mean compliance as well as its variability. In Figure 2.20(a), we analyze the impact of mean compliance changing from 10% to 90%, whereas the range of the uniform distribution is kept constant at 20%; for example, if the mean compliance is 10%, the compliance uniformly belongs to the range $[0, 20\%]$. Similarly, in Figure 2.20(b), we test the changes in the variability by varying the range size between 0% and 80%, while keeping the mean compliance constant at 50%; for example, if the range size is 80%, the compliance uniformly belongs to the range $[10\%, 90\%]$. To obtain each data point, we use 300 training and 300 testing realizations. To compute the expected cost of our solution for a testing realization, after HWP assigns calls to today, we identify the number of active calls in each period and use our pre-computed values $\varpi_{r,\mathcal{R}}$ to determine the expected cost.

The error values in Figure 2.20 are computed by comparing the expected cost of our solution with the lower bound. Because no assumption is made while computing our lower bound, the reported errors in Figure 2.20 include all errors in our approximation. We observe, under stochastic compliance, the total error of our extended approximation remains small and the performance of our approximation is robust when mean and/or variability of compliance

change. As one would expect, we observe a modest increase in the error when the mean compliance decreases and/or when the variability increases. We also observe in Figure 2.20(a) a large error for the case of mean compliance of 10%, which is an unrealistic case in our application because DLCCs have significantly high compliance rates. Nevertheless, our approximation does not perform well when the mean compliance is significantly small.

Note as compliance variability increases, the odds of achieving the threshold decreases. In our application, the compliance variability is small and the threshold is not strictly enforced, and hence, our above extension is appropriate for our application. Ahead we discuss extending our approximation if achieving the threshold (with a given probability , e.g., 95%) is strictly enforced by the utility.

Insights on why our generalization works. We provide an intuition on the excellent performance of our extension for the case of uncertain compliance. Recall the expected cost in period t is

$$\varpi_{r,k} := \mathbf{E}_{\chi_1,\ldots,\chi_k}\left[f\left(r - \sum_{i=1}^{k}\chi_i\right)\right],$$

whereas we schedule calls using an approximate expected cost

$$\hat{\varpi}_{r,k} := f\left(r - \sum_{i=1}^{k}\mathbf{E}[\chi_i]\right).$$

Because f is a smooth function and the variability of χ_i is low (consistent with our practical application), the two values $\varpi_{r,k}$ and $\hat{\varpi}_{r,k}$ are very close. Furthermore, as $\varpi_{r,k} \geq \hat{\varpi}_{r,k}$, we are always underestimating the cost. Hence, using $\hat{\varpi}_{r,k}$ to schedule calls is a good approximation. After determining the schedule of calls, we use the exact values $\varpi_{r,k}$ to compute the expected cost.

Achieving the threshold with a given probability. Our generalized approximation achieves the threshold in expectation, whereas an actual realization may slightly exceed the threshold. This approach is appropriate in our application because in practice, achieving the threshold

is desirable but strictly enforcing the threshold is not necessary. We briefly discuss the implications of strictly enforcing the threshold below. Assume a utility requires achieving the threshold with a given probability a at least. Assume χ_is are i.i.d. For each period t, we can compute the number of active calls needed to reduce the consumption to the threshold with probability a. Thus, given a daily consumption profile and a threshold, with certainty, we determine the number of active calls needed in each period of the day. Given this information, the rest of our approach extends as we explained above. The only difference between our former generalization and this one is that- if the compliance variability is high, in the latter case, we may need to assign additional calls to a period so that we can achieve the threshold with probability a; however, the former case schedules calls solely based on the expected compliance rate.

2.8 Conclusion

The high cost and long lead time of investments in power systems combined with the lack of efficient storage technologies makes reducing peak loads an important yet challenging task. The importance of managing peak loads is growing with rapid changes in the climate. Climate change is resulting in extreme weather that in turn is causing record-setting peak-energy-consumption levels.[9] In this research work, we study direct load control contracts (DLCCs), a class of incentive-based demand-response programs that enable utility firms to reduce the load during peak periods. DLCCs are becoming more popular among utilities. Efficient implementation of these contracts is not a trivial task, because the energy-consumption pattern varies across days, and monthly and annual limits are placed on the number of times and hours that customers can be called to reduce their load.

Experiments with data from CAISO verify that using our solution approach for managing

[9] https://www.cbsnews.com/sanfrancisco/news/california-rolling-blackouts-power-grid-heat-wave-hot-weather/ (retrieved on September 12, 2022).

peak load is highly beneficial. Moreover, a large utility firm in California that used our optimization framework achieved a significant reduction in cost. Our model also provided insights for managers about the effects of changing the features of DLCCs. For example, we observed that our model maintains a low variability in the number of calls and hours across the groups, and it calls customers for longer hours if necessary, but the calls are made less frequently. Because our modeling framework captures the most salient features of DLCCs, and the solution scheme is aligned with managerial considerations, our analytical framework can be implemented in other utility firms that use DLCCs, perhaps with some customization. For example, if a utility allows daily/weekly thresholds, our approach can be modified to incorporate this additional feature. For this scenario, we can also theoretically analyze the worst-case error of using flat-tops solutions when calls are at most one hour.

The DLCCs studied in this chapter contractually indicate that load shifting is not allowed. Our numerical analysis indicated the load-shifting is not likely to be a challenge in our application, because our approximation could be used if load shifting does not create new peaks. Future research needs to perform a more thorough analysis of these variations of DLCCs by modeling the stochastic behavior of customers in response to calls.

Finally, the DLCC studied in this work captures most of the features of the DLCC contracts offered by many utility companies. However, more research is needed to address the contracts that allow non-equal group size and dynamic enrolment in these programs.

2.9 Appendix A: Mathematical Model for the Deterministic Version of P

In the deterministic special case of P, denoted by DDP, we know (with certainty) the consumption rate at any time over the contract horizon. In this deterministic setting, let r_{mdt} denote the energy consumption during period t of day d in month m. Let binary variable $x_{mdgt} \in \{0,1\}$ be 1 if group g is called in period t of day d in month m, and 0 otherwise. Quantity $\sum_{g=1}^{G} x_{mdgt}$ is the number of groups that are on call in period t of day d in month m.

We assume each hour of a call reduces the load by one unit (e.g., 1 GWh). This assumption is without loss of generality because if, instead, each hour of a call reduces the load by $\bar{r}$ GWh, r_{mdt} would be measured as $\bar{r}$ GWh. Therefore, after assigning calls to groups, the energy consumption in period t of day d in month m becomes $r_{mdt} - \sum_{g=1}^{G} x_{mdgt}$. We also assume, for ease of presentation, r_{mdt} is sufficiently large such that $r_{mdt} - \sum_{g=1}^{G} x_{mdgt} \geq 0$, for all m, d, and t.

Let $f_{mdt}(r_{mdt} - \sum_{g=1}^{G} x_{mdgt})$ denote the cost function in period t of day d in month m if the consumption (before reduction) is r_{mdt} and the number of groups on call is $\sum_{g=1}^{G} x_{mdgt}$. In addition, let binary variable $y_{mdgt} \in \{0,1\}$ be 1 if a call to group g starts in period t of day d in month m, and 0 otherwise. DDP is formulated as

$$\text{(DDP)} \qquad \min \quad \sum_{m=1}^{\mathcal{M}} \sum_{d=1}^{\mathcal{D}_m} \sum_{t=1}^{T} f_{mdt}\Big(r_{mdt} - \sum_{g=1}^{G} x_{mdgt}\Big) \tag{2.8a}$$

$$\text{s.t.} \quad \sum_{m=1}^{\mathcal{M}} \sum_{d=1}^{\mathcal{D}_m} \sum_{t=1}^{T} y_{mdgt} \leq K, \quad \forall g, \tag{2.8b}$$

$$\sum_{m=1}^{\mathcal{M}} \sum_{d=1}^{\mathcal{D}_m} \sum_{t=1}^{T} x_{mdgt} \leq H, \quad \forall g, \tag{2.8c}$$

$$\sum_{d=1}^{\mathcal{D}_m} \sum_{t=1}^{T} y_{mdgt} \leq \breve{k}_m, \quad \forall m,\ g, \tag{2.8d}$$

$$\sum_{d=1}^{\mathcal{D}_m}\sum_{t=1}^{T} x_{mdgt} \leq \breve{h}_m, \quad \forall m,\ g, \tag{2.8e}$$

$$\sum_{t=1}^{T} x_{mdgt} \leq L, \quad \forall m,\ g,\ d = 1,\ldots,\mathcal{D}_m, \tag{2.8f}$$

$$y_{mdgt} \geq x_{mdgt} - x_{mdg,t-1}, \quad \forall m,\ g,\ t \geq 2,\ d = 1,\ldots,\mathcal{D}_m, \tag{2.8g}$$

$$y_{mdg1} \geq x_{mdg1}, \quad \forall m,\ g,\ d = 1,\ldots,\mathcal{D}_m, \tag{2.8h}$$

$$\sum_{t=1}^{T} y_{mdgt} \leq 1, \quad \forall m,\ g,\ d = 1,\ldots,\mathcal{D}_m, \tag{2.8i}$$

$$x_{mdgt}, y_{mdgt} \in \{0,1\}, \quad \forall m,\ g,\ t,\ d = 1,\ldots,\mathcal{D}_m. \tag{2.8j}$$

The objective function (2.8a) minimizes the total energy-consumption cost throughout the planning horizon. Constraints (2.8b)-(2.8e) impose the yearly and monthly limits on the total number of calls and hours assigned to each group. Constraint (2.8f) ensures the duration of each call does not exceed L. Constraints (2.8g)-(2.8i) guarantee the continuity of calls and that each group is called at most once in one day.

2.10 Appendix B: Proofs

2.10.1 Proof of Proposition 1

Proof of part (a). Consider the *2-Partition* problem [GJ79] with $2n$ elements $A = \{a_1,\ldots,a_{2n}\}$, such that $a_i \in \mathbb{Z}^+$, for all i, and $\sum_{i=1}^{2n} a_i = 2B$. Can A be divided into two disjoint subsets of size n, such that the sum of the elements in each subset equals B? Consider a specific instance of the deterministic version of P, denoted by DDP: $G = 2, \mathcal{M} = 1, K = \breve{k}_1 = n, H = \breve{h}_1 = B, \mathcal{D}_1 = 2n$, $T \geq \max_i\{a_i\}$, and $L \geq \max_i\{a_i\}$. The i^{th} day has an ECP of length a_i and height 1 such that if a call of length a_i is assigned to a group on the i^{th} day, a saving of a_i will be achieved, whereas if the group cannot receive the full-length call because of its contractual limits, the saving will be strictly less than a_i. Note the total load over the horizon

is equal to $2B$.

Suppose the 2-partition problem has a feasible solution that divides A into disjoint subsets A_1 and A_2 of size n. Then, by assigning the ECPs in A_1 (i.e., calls of length $a_i \in A_1$) to one of the groups and the ECPs in A_2 (i.e., calls of length $a'_i \in A_2$) to the other group, we can achieve a total saving of $2B$. If the 2-partition problem has no feasible solution, the sum of the elements in either A_1 or A_2 exceeds B. Without loss of generality, assume the sum of the elements in A_1 (A_2) is larger (smaller) than B. Given that each group can be called up to n times for a total duration of up to B hours, the saving from A_1 is B and the saving from A_2 is strictly less than B. Thus, the total saving is strictly less than $2B$.

Proof of part (b). We use the *Numerical 3-Dimensional Matching* problem (NMP), which is known to be strongly NP-complete [GJ79]. Consider disjoint sets $C_1 = \{a_1, \ldots, a_n\}, C_2 = \{a_{n+1}, \ldots, a_{2n}\}$, and $C_3 = \{a_{2n+1}, \ldots, a_{3n}\}$, each containing n elements of sizes $a_i \in \mathbb{Z}^+$ for $1 \leq i \leq 3n$, and let $B \in \mathbb{Z}^+$. Can $C_1 \cup C_2 \cup C_3$ be partitioned into n disjoint subsets $A_1, A_2, \ldots, A_n$, such that each A_i contains exactly one element from each of C_1, C_2, and C_3, and $\sum_{a_i \in A_j} a_i = B$ for $1 \leq j \leq n$?

Consider the following instance of DDP: $G = n$, $\mathcal{M} = 3$, $\mathcal{D}_m = n$, $\breve{k}_m = 1$, $\breve{h}_m \geq \max_i\{a_i\}$, for all m, $K = 3$, $H = B$, $T \geq \max_i\{a_i\}$, and $L \geq \max_i\{a_i\}$. The energy consumption over the three month horizon is illustrated below.

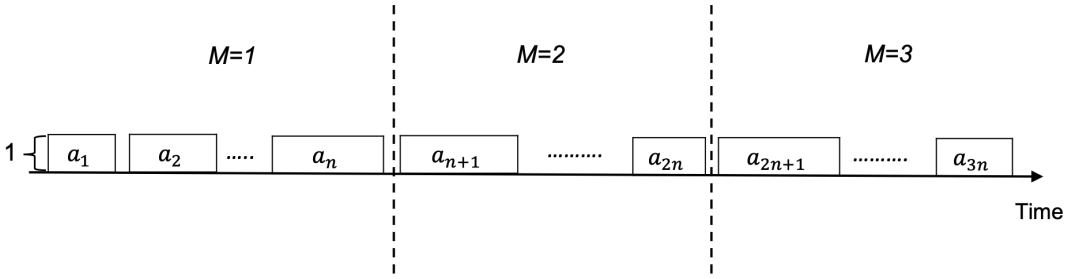

Suppose NMP has a feasible solution. Then, by assigning the ECPs in A_j (i.e., the 3 calls, 1 call per month) to group $1 \leq j \leq n$, we can achieve a total saving of nB. If NMP does not have a feasible solution, a subset exists with the sum of elements exceeding B. Because each group can be called at most once per month and for a total of B hours over the three

months, the total saving will be strictly less than nB.

Proof of part (c). Consider the *3-Partition* problem [GJ79] with the set of $3n$ elements $A = \{a_1, \ldots, a_{3n}\}$, such that $a_i \in \mathbb{Z}^+$, $\frac{B}{4} < a_i < \frac{B}{2}$, for all i, and $\sum_{i=1}^{3n} a_i = nB$. Consider an instance of DDP where $G = 3n$, group i has one remaining call and remaining hours a_i, and $L \geq \max_i\{a_i\}$. Consider the following ECP for the last day, when a total load of nB is distributed over B hours, and the load is zero for the rest of the day.

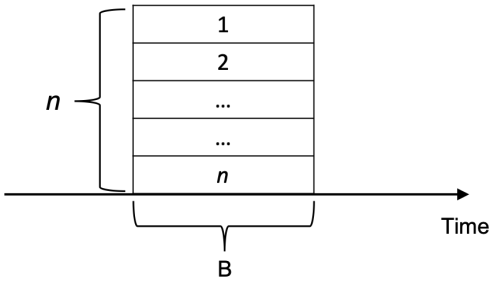

Suppose the 3-partition problem has a feasible solution and A can be divided into n disjoint subsets $A_1, \ldots, A_n$ of size 3, such that the sum of the elements in each subset is B. Then, we assign the j^{th} layer of the ECP to the three groups for which $a_i \in A_j$, for all $1 \leq j \leq n$ and achieve a total saving of nB. If the 3-partition problem does not have a feasible solution, subsets A_j and $A_{j'}$ exist such that the sum of their elements is larger and smaller than B, respectively. Because each group has one remaining call, the saving from A_j will be B, whereas the saving from $A_{j'}$ will be less than B; thus, the total saving across all subsets will be strictly less than nB.

2.10.2 Proof of Theorem 1

At each iteration of the *while*-loop, Algorithm 1 creates a call, denoted by $\bar{\eta}$, with starting time t^s and end time t^e. Consider the first iteration of the *while*-loop. If the energy consumption never exceeds $\mathcal{Z}$, t^s does not exist, and no call is created in the *while*-loop. Otherwise, a call is created in the first iteration of the *while*-loop. Let us denote this call by $\bar{\eta}$, its starting time by t^s, and its ending time by t^e. We need to show an optimal solution for DLP exists

that satisfies the following condition: j' exists such that $x_{j't} = 1$, for all $t \in \{t^s, t^s+1, \ldots, t^e\}$, $x_{j't} = 0$, for all $t \notin \{t^s, t^s+1, \ldots, t^e\}$, $y_{j't^s} = 1$, and $y_{j't} = 0$, for all $t \neq t^s$. After we prove this result, we fix the values of $x_{j't}$s and $y_{j't}$s and solve the reduced problem. The reduced problem has the same structure with $J-1$ potential calls, and the energy consumption has been reduced by one unit during hours t^s to t^e. By repeating this logic, the *while*-loop finds the minimum number of calls to eliminate the area above $\mathcal{Z}$.

Consider the first iteration of the *while*-loop and assume t^s exists such that $t^s = \min\{t|\ 1 \leq t \leq T,\ \bar{r}_t > \mathcal{Z}\}$. Because of constraint (2.3b), in an optimal solution of DLP, j' exists such that $x_{j't^s} = 1$. The algorithm creates a call with starting time t^s and end time t^e such that $t^e = \min\{t^s + L - 1, \min\{t|t^s \leq t \leq T,\ \bar{r}_t > \bar{r}_{t+1}\}\}$. We aim to show an optimal solution for DLP exists such that $x_{j't} = 1$, for all $t \in \{t^s, t^s+1, \ldots, t^e\}$, and $x_{j't} = 0$, for all $t \notin \{t^s, t^s+1, \ldots, t^e\}$. Note either $t^s = 1$ or $x_{j't} = 0$, for all $t \in \{1, 2, \ldots, t^s - 1\}$ (because of constraint (2.3b)). Therefore, we need to prove $x_{j't} = 1$, for all $t \in \{t^s, t^s+1, \ldots, t^e\}$, and $x_{j't} = 0$, for all $t \in \{t^e+1, t^e+2, \ldots, T\}$. Note if $t^e = T$, we only need to prove $x_{j't} = 1$, for all $t \in \{t^s, t^s+1, \ldots, t^e\}$; hence, without loss of generality, assume $t^e < T$. Moreover, it is sufficient to show $x_{j't} = 1$, for all $t \in \{t^s, t^s+1, \ldots, t^e\}$, and $x_{j't^e+1} = 0$ (because of constraints (2.3d)-(2.3g)).

We first prove an optimal solution for DLP exists such that $x_{j't} = 1$, for all $t \in \{t^s, t^s+1, \ldots, t^e\}$. Suppose to the contrary that in any optimal solution for DLP, t' exists such that $t' := \min\{t|\ t^s + 1 \leq t \leq t^e,\ x_{j't} = 0\}$. Note $\mathcal{Z} < \bar{r}_{t'-1} \leq \bar{r}_{t'}$, and hence, $\sum_j x_{j,t'-1} \leq \sum_j x_{j,t'}$ (because of constraint (2.3b)). Therefore, j'' exists such that $x_{j'',t'-1} = 0$ and $x_{j'',t'} = 1$. Moreover, $y_{j'',t'} = 1$. This observation means that two calls are made: one is active during t^s to $t'-1$ and the other is active during t' to t'', for some $t'' \geq t'$. We change this optimal solution as follows: let the first call be active during t^s to t' and the second call be active during $t'+1$ to t'' (if $t' = t''$ the second call is eliminated meaning the objective value of DLP decreases by 1, and if $t' < t''$, the objective value remains the same). The duration of the first

call is increased by 1 and becomes $t' - t^s + 1$. This increase satisfies constraint (2.3c) because

$$t' - t^s + 1 \leq t^e - t^s + 1 \leq (t^s + L - 1) - t^s + 1 \leq L.$$

Moreover, note the duration of the second call is reduced and it automatically satisfies constraint (2.3c). Thus, the modified solution is feasible for DLP and its objective value is at least as good as that of the original solution. Therefore, the modified solution is also optimal and t' has increased by 1. By repeating this procedure while t' exists satisfying $t' := \min\{t|\ t^s+1 \leq t \leq t^e,\ x_{j't} = 0\}$, we create an optimal solution such that no such t' exists. Hence, an optimal solution for DLP exists such that $x_{j't} = 1$, for all $t \in \{t^s, t^s + 1, \ldots, t^e\}$.

Next, we aim to prove an optimal solution for DLP exists such that $x_{j't} = 1$, for all $t \in \{t^s, t^s + 1, \ldots, t^e\}$, and $x_{j't^e+1} = 0$. Consider the following two cases.

Case 1: $t^s + L - 1 > \min\{t|t^s \leq t \leq T,\ \bar{r}_t > \bar{r}_{t+1}\}$. Then, $t^e = \min\{t|t^s \leq t \leq T,\ \bar{r}_t > \bar{r}_{t+1}\}$ and $t^e < t^s + L - 1$. Suppose to the contrary that in any optimal solution that satisfies $x_{j't} = 1$, for all $t \in \{t^s, t^s + 1, \ldots, t^e\}$, for some j', we have $x_{j't^e+1} = 1$. Therefore, consider an optimal solution such that j' exists for which $x_{j't} = 1$, for all $t \in \{t^s, t^s + 1, \ldots, t^e, t^e + 1\}$. Let t' be such that $x_{j't'} = 1$ and $x_{j't'+1} = 0$. Thus, $t' \geq t^e + 1$ and $x_{j',t} = 1$, for $t \in \{t^s, t^s + 1, \ldots, t'\}$, and $x_{j',t} = 0$, for $t \notin \{t^s, t^s + 1, \ldots, t'\}$. In other words, we have a call that is active during t^s to t'. Note

$$\begin{aligned} & \mathcal{Z} < \bar{r}_{t^s} \leq \bar{r}_{t^s+1} \leq \cdots \leq \bar{r}_{t^e} > \bar{r}_{t^e+1} \\ \Rightarrow \quad & \sum_j x_{j,t^s} \leq \sum_j x_{j,t^s+1} \leq \cdots \leq \sum_j x_{j,t^e} > \sum_j x_{j,t^e+1}, \end{aligned}$$

because of constraint (2.3b). Therefore, j'' and $t'' \in \{t^s + 1, \ldots, t^e\}$ exist such that $x_{j'',t} = 1$, for $t \in \{t'', t'' + 1, \ldots, t^e\}$, and $x_{j'',t} = 0$, for $t \notin \{t'', t'' + 1, \ldots, t^e\}$. We modify this optimal solution as follows: $x_{j',t} = 0$, for $t \in \{t^e + 1, \ldots, t'\}$, and $x_{j'',t} = 1$, for $t \in \{t^e + 1, \ldots, t'\}$ (we do not change the values of other variables). This modified solution is feasible for DLP and

its objective value is the same. Therefore, we have created an optimal solution satisfying $x_{j't} = 1$, for all $t \in \{t^s, t^s + 1, \ldots, t^e\}$, and $x_{j't^e+1} = 0$.

Case 2: $t^s + L - 1 \leq \min\{t | t^s \leq t \leq T,\ \bar{r}_t > \bar{r}_{t+1}\}$. Then, $t^e = t^s + L - 1$. In this case, $x_{j't^e+1} = 0$ because of constraint (2.3c).

So far, we have proven Algorithm 1 finds an optimal solution for DLP. Next, we argue this optimal solution is found in time complexity $\mathcal{O}(JT)$. The *while*-loop iterates at most J times. At each iteration of the *while*-loop, t^s is found (or its nonexistence is verified) in $\mathcal{O}(T)$. In line 3, t^e is also found in $\mathcal{O}(T)$. Line 5 requires at most L operations ($\leq \mathcal{O}(T)$). Therefore, each iteration of the *while*-loop is performed in time complexity of $\mathcal{O}(T)$. Thus, the computational time complexity of Algorithm 1 is $\mathcal{O}(JT)$, and hence, the proof is complete.

2.10.3 Proof of Theorem 2

We first provide an intuitive explanation for this error bound. To prove Theorem 2, we consider an optimal solution for the aggregate problem and create a feasible solution for the disaggregate problem. We show the portion of the aggregate solution that corresponds to a month can be disaggregated such that the total hours assigned to any two groups differ by at most L hours, and thereby, in each month, for each group, at most δL error may be created. This step leads to a worst-case error of $\delta G\mathcal{M}L$, corresponding to the first part of the error bound in Theorem 2. The second part of this bound is created while combining the monthly solutions. We rotationally assign the monthly solutions to the groups. After this step, the total hours assigned to any two groups differ by at most $\mathcal{M}L$ hours, and hence, we create at most $\delta G\mathcal{M}L$ error. Note, however, combining the monthly solutions may create no (little) error if the annual limit H imposes no (little) further restriction on top of the monthly limits. Obviously, if the annual limit H is relaxed, our rotational assignment creates no further error. Therefore, the error we create in the second step depends on the restrictiveness of the annual limit H and cannot be more than $\delta\left(\left(\sum_m \breve{h}_m\right) - H\right)$ for each group. Thus, the total error of the second step is bounded by the minimum of the two quantities, $\delta G\mathcal{M}L$ and

$\delta G\left(\left(\sum_m \breve{h}_m\right) - H\right)$.

Let DDP and ADDP denote the deterministic versions of P and AMP without restricting solutions to threshold-type policies. For ease of presentation, let v_{DDP} and v_{ADDP} respectively denote the optimal values of DDP and ADDP (these quantities are denoted by v' and v'' in the main manuscript). Note ADDP is a relaxation of DDP, because a feasible solution of DDP is always feasible for ADDP but the converse may not be true. Thus, $v_{\text{DDP}} \geq v_{\text{ADDP}}$, which implies 0 is a lower bound on $(v_{\text{DDP}} - v_{\text{ADDP}})$. Let $\mathcal{C}^A$ denote the set of calls in the optimal solution of ADDP. For call $\eta \in \mathcal{C}^A$, let $m(\eta),\ d(\eta),\ t(\eta)$, and $l(\eta)$ denote the month, day, starting time, and duration of η. Our proof is based on showing a feasible solution for DDP exists, denoted by $\mathcal{C}$, such that the objective value of $\mathcal{C}$, denoted by $v(\mathcal{C})$, satisfies

$$v(\mathcal{C}) - v_{\text{ADDP}} \leq \delta G\left(\mathcal{M}L + \min\left\{\mathcal{M}L,\ \left(\sum_{m=1}^{\mathcal{M}} \breve{h}_m\right) - H\right\}\right). \qquad (\dagger)$$

After showing this result, and because $v_{\text{ADDP}} \leq v_{\text{DDP}} \leq v(\mathcal{C})$, we have

$$v_{\text{DDP}} - v_{\text{ADDP}} \leq v(\mathcal{C}) - v_{\text{ADDP}} \leq \delta G\left(\mathcal{M}L + \min\left\{\mathcal{M}L,\ \left(\sum_{m=1}^{\mathcal{M}} \breve{h}_m\right) - H\right\}\right),$$

which completes the proof of Theorem 2.

Therefore, we need to prove $\mathcal{C}$ exists that satisfies (†). We present a procedure, given in Algorithm 2, referred to as FDS (which stands for *feasible disaggregate solution*), which creates a feasible solution for DDP.

In fact, given $\mathcal{C}^A$, we aim to determine the recipient group for each call $\eta \in \mathcal{C}^A$ (the duration of some of the calls may be reduced). Let $g(\eta)$ denote the group that receives call $\eta \in \mathcal{C}^A$. The output of procedure FDS is a set of calls $\mathcal{C}$, where for each call $\eta \in \mathcal{C}$ the values of $m(\eta),\ d(\eta),\ t(\eta)$, $l(\eta)$, and $g(\eta)$ are known. In line 1 of procedure FDS, $\bar{\mathcal{C}}_m$ contains all calls in month m, and we aim to assign these calls to groups. These calls are sorted in line 3. The first G calls are referred to as class 1 (CLS(η)=1), the second G calls are referred to as

Algorithm 2: Procedure FDS.

Input: $\mathcal{C}^A$; ▷ `Attributes` $m(\eta)$, $d(\eta)$, $t(\eta)$, `and` $l(\eta)$ `are known` $\forall \eta \in \mathcal{C}^A$

Output: $\mathcal{C}$; ▷ `Attributes` $m(\eta)$, $d(\eta)$, $t(\eta)$, $l(\eta)$, `and` $g(\eta)$ `are known` $\forall \eta \in \mathcal{C}$

Define $\mathcal{C} := \{\}$;

Define $\bar{\mathcal{C}}_m := \{\eta \in \mathcal{C}^A \,|\, m(\eta) = m\}$, for all $m = 1, \ldots, \mathcal{M}$.

for $m = 1, 2, \ldots, \mathcal{M}$ **do**

- Sort the calls in $\bar{\mathcal{C}}_m$ in a non-increasing order of their durations.
- Let $\text{ORD}(\eta) \in \{1, \ldots, |\bar{\mathcal{C}}_m|\}$ denote the order of η in the sorted list, for all $\eta \in \bar{\mathcal{C}}_m$.
- Define $\text{CLS}(\eta) := \lceil \text{ORD}(\eta)/G \rceil$, for all $\eta \in \bar{\mathcal{C}}_m$.
- Partition $\bar{\mathcal{C}}_m$ into $\bar{\bar{\mathcal{C}}}_{m1}, \ldots, \bar{\bar{\mathcal{C}}}_{mG}$ satisfying
  - (i) $|\bar{\bar{\mathcal{C}}}_{mg}| \geq |\bar{\bar{\mathcal{C}}}_{m(g+1)}|$, for all $g = 1, \ldots, G-1$, and
  - (ii) $d(\eta') \neq d(\eta'')$ and $\text{CLS}(\eta') \neq \text{CLS}(\eta'')$, for all g, for all $\eta',\ \eta'' \in \bar{\bar{\mathcal{C}}}_{mg}$.
- **for** $g = 1, 2, \ldots, G$ **do**
  - While $\sum_{\eta \in \bar{\bar{\mathcal{C}}}_{mg}} l(\eta) > \breve{h}_m$, reduce $l(\eta)$ for some $\eta \in \bar{\bar{\mathcal{C}}}_{mg}$.
  - Let $g' := \arg\min_{g''=1,\ldots,G} \{\sum_{\eta \in \mathcal{C}} \mathbb{I}(g(\eta) = g'')\}$.
  - Set $g(\eta) = g'$, for all $\eta \in \bar{\bar{\mathcal{C}}}_{mg}$.
  - Update $\mathcal{C} := \mathcal{C} \cup \bar{\bar{\mathcal{C}}}_{mg}$.

for $g = 1, 2, \ldots, G$ **do**

- While $\sum_{\eta \in \mathcal{C}} l(\eta) \mathbb{I}(g(\eta) = g) > H$, reduce $l(\eta)$ for some η such that $\eta \in \mathcal{C}$, $g(\eta) = g$.

class 2 (CLS(η)=2), and so on (lines 4 and 5). In line 6, $\bar{\mathcal{C}}_m$ is partitioned into G mutually exclusive and collectively exhaustive subsets $\bar{\bar{\mathcal{C}}}_{m1}, \ldots, \bar{\bar{\mathcal{C}}}_{mG}$ such that each $\bar{\bar{\mathcal{C}}}_{mg}$ contains at most one call from each class and at most one call from each day. In addition, we require that the subsets $\bar{\bar{\mathcal{C}}}_{mg}$s are sorted in a non-increasing order of their sizes. In lines 7-11, we iteratively assign the largest subset to the group with the least assigned calls. We also reduce the durations of some of the calls in line 8 so that the total hours assigned to a group in a month is less than or equal to $\breve{h}_m$. Finally, in lines 13-15, we reduce the durations of some of the calls so that the total hours assigned to each group is less than or equal to H. Lemma 1 establishes feasibility of procedure FDS.

Lemma 1. *Given an optimal solution of ADDP, procedure FDS creates a feasible solution for DDP.*

Proof. We need to prove the produced set of calls $\mathcal{C}$ is a feasible solution for DDP. Note

ADDP has the following constraint:

$$\sum_{g=1}^{G}\sum_{d=1}^{\mathcal{D}_m}\sum_{t=1}^{T} y_{mdgt} \leq G\breve{k}_m, \quad \forall m.$$

Because of this constraint, we have $|\bar{\mathcal{C}}_m| \leq G\breve{k}_m$, for all m. At iteration m, after creating classes, some classes have exactly G calls and at most one class has less than G calls. Consider a class with exactly G calls. Note we partition $\bar{\mathcal{C}}_m$ into G subsets and each subset can have at most one call from a class; therefore, each subset $\bar{\bar{\mathcal{C}}}_{mg}$ contains exactly one call from a class that has exactly G calls. Thus, if all classes have exactly G calls, all subsets $\bar{\bar{\mathcal{C}}}_{mg}$s will contain an equal number of calls. Now, consider a class with less than G calls. Each subset receives at most one call from this class. Therefore, we have

$$\left||\bar{\bar{\mathcal{C}}}_{mg'}| - |\bar{\bar{\mathcal{C}}}_{mg''}|\right| \leq 1, \quad \forall m,\ g',\ g''. \tag{2.9}$$

Next, we use a contradiction to show $|\bar{\bar{\mathcal{C}}}_{mg}| \leq \breve{k}_m$, for all m and g. Suppose m' and g' exist such that $|\bar{\bar{\mathcal{C}}}_{m'g'}| > \breve{k}_{m'}$. Thus, because of inequality (2.9), we must have $|\bar{\bar{\mathcal{C}}}_{m'g}| \geq \breve{k}_{m'}$, for all $g \neq g'$. This observation implies that $|\bar{\mathcal{C}}_{m'}| = \sum_{g=1}^{G} |\bar{\bar{\mathcal{C}}}_{m'g}| > G\breve{k}_{m'}$, which is a contradiction because $|\bar{\mathcal{C}}_{m'}| \leq G\breve{k}_{m'}$. Therefore, we proved $|\bar{\bar{\mathcal{C}}}_{mg}| \leq \breve{k}_m$, for all $m,\ g$. A subset $\bar{\bar{\mathcal{C}}}_{mg}$ contains some calls in month m that are assigned to a group in line 10. Therefore, procedure FDS satisfies constraint (2.8d). Moreover, in line 8, we reduce the durations of some of the calls so that the sum of durations of calls in $\bar{\bar{\mathcal{C}}}_{mg}$ is less than or equal to $\breve{h}_m$. Thus, procedure FDS satisfies constraint (2.8e).

Next, we show procedure FDS also satisfies constraints (2.8b) and (2.8c). Using an induction, we show that at the end of iteration m of the *for*-loop (lines 2-12), set $\mathcal{C}$ satisfies

$$\left|\left(\sum_{\eta\in\mathcal{C}} \mathbb{I}(g(\eta) = g')\right) - \left(\sum_{\eta\in\mathcal{C}} \mathbb{I}(g(\eta) = g'')\right)\right| \leq 1, \quad \forall g',\ g'', \tag{2.10}$$

meaning the number of calls assigned to different groups differs by at most one. At the end of iteration $m = 1$, inequality (2.10) is satisfied because of inequality (2.9). Assume inequality (2.10) is satisfied at the end of iteration m. Thus, K' exists such that some groups have been assigned K' calls and some groups have been assigned $K' + 1$ calls during months 1 to m. Let $\mathcal{G}'$ denote the set of groups that have been assigned K' calls and $\mathcal{G}''$ denote the set of groups that have been assigned $K' + 1$ calls. In month $m + 1$, we want to assign $\bar{\bar{\mathcal{C}}}_{m+1,1}, \ldots, \bar{\bar{\mathcal{C}}}_{m+1,G}$ to groups. Note g' and K'' exists such that $\bar{\bar{\mathcal{C}}}_{m+1,g}$ contains $K'' + 1$ calls, for all $g = 1, \ldots, g'$, and $\bar{\bar{\mathcal{C}}}_{m+1,g}$ contains K'' calls, for all $g = g' + 1, \ldots, G$. Exactly one of the following two cases happens.

Case 1: $|\mathcal{G}'| \geq g'$. Then, let $\mathcal{G}'_{sub} \subseteq \mathcal{G}'$ such that $|\mathcal{G}'_{sub}| = g'$. Assume without loss of generality that $\bar{\bar{\mathcal{C}}}_{m+1,1}, \ldots, \bar{\bar{\mathcal{C}}}_{m+1,g'}$ are assigned to groups in $\mathcal{G}'_{sub}$. At the end of iteration $m+1$, we have: $\sum_{\eta \in \mathcal{C}} \mathbb{I}(g(\eta) = g) = K' + K'' + 1$, for all $g \in \mathcal{G}'_{sub} \cup \mathcal{G}''$, and $\sum_{\eta \in \mathcal{C}} \mathbb{I}(g(\eta) = g) = K' + K''$, for all $g \notin \mathcal{G}'_{sub} \cup \mathcal{G}''$.

Case 2: $|\mathcal{G}'| < g'$. Then, let $\mathcal{G}''_{sub} \subseteq \mathcal{G}''$ such that $|\mathcal{G}''_{sub}| = G - g'$. Assume without loss of generality that $\bar{\bar{\mathcal{C}}}_{m+1,g'+1}, \ldots, \bar{\bar{\mathcal{C}}}_{m+1,G}$ are assigned to groups in $\mathcal{G}''_{sub}$. At the end of iteration $m + 1$, we have: $\sum_{\eta \in \mathcal{C}} \mathbb{I}(g(\eta) = g) = K' + K'' + 1$, for all $g \in \mathcal{G}' \cup \mathcal{G}''_{sub}$, and $\sum_{\eta \in \mathcal{C}} \mathbb{I}(g(\eta) = g) = K' + K'' + 2$, for all $g \notin \mathcal{G}' \cup \mathcal{G}''_{sub}$.

Therefore, in either cases, inequality (2.10) is satisfied at the end of iteration $m + 1$. Moreover, inequality (2.10) is satisfied at the end of iteration $m = \mathcal{M}$.

Next, we prove that at the end of iteration $m = \mathcal{M}$, $\sum_{\eta \in \mathcal{C}} \mathbb{I}(g(\eta) = g) \leq K$, for all g, using a contradiction. Suppose g' exists such that $\sum_{\eta \in \mathcal{C}} \mathbb{I}(g(\eta) = g') > K$. Thus, because of inequality (2.10), we must have $\sum_{\eta \in \mathcal{C}} \mathbb{I}(g(\eta) = g) \geq K$, for all $g \neq g'$. This observation implies $|\mathcal{C}^A| = \sum_{g=1}^{G} \sum_{\eta \in \mathcal{C}} \mathbb{I}(g(\eta) = g) > GK$, which is a contradiction because $|\mathcal{C}^A| \leq GK$ due to the following constraint in ADDP:

$$\sum_{g=1}^{G} \sum_{m=1}^{\mathcal{M}} \sum_{d=1}^{\mathcal{D}_m} \sum_{t=1}^{T} y_{mdgt} \leq GK.$$

Therefore, procedure FDS satisfies constraint (2.8b). Moreover, in line 14, we reduce the durations of some of the calls so that the sum of durations of calls that are assigned to each group is less than or equal to H. Thus, procedure FDS satisfies constraint (2.8c). It follows that the produced set of calls $\mathcal{C}$ is a feasible solution for DDP, and hence, the proof is complete. □

So far, we have proven $\mathcal{C}$ is a feasible solution for DDP. Next, we show $\mathcal{C}$ satisfies (†). The values of v_{ADDP} and $v(\mathcal{C})$ are obtained using the following equations:

$$v_{\text{ADDP}} = \sum_{m=1}^{\mathcal{M}} \sum_{d=1}^{\mathcal{D}_m} \sum_{t=1}^{T} f\Big(r_{mdt} - \sum_{\eta \in \mathcal{C}^A:\ m(\eta)=m,\ d(\eta)=d} \mathbb{I}\Big(t(\eta) \le t < t(\eta) + l(\eta)\Big)\Big), \quad (2.11)$$

$$v(\mathcal{C}) = \sum_{m=1}^{\mathcal{M}} \sum_{d=1}^{\mathcal{D}_m} \sum_{t=1}^{T} f\Big(r_{mdt} - \sum_{\eta \in \mathcal{C}:\ m(\eta)=m,\ d(\eta)=d} \mathbb{I}\Big(t(\eta) \le t < t(\eta) + l(\eta)\Big)\Big). \quad (2.12)$$

Using equation (2.12), we have

$$v(\mathcal{C}) = \sum_{m=1}^{\mathcal{M}} \sum_{d=1}^{\mathcal{D}_m} \sum_{t=1}^{T} f\Big(r_{mdt} - \sum_{\eta \in \mathcal{C}:\ m(\eta)=m,\ d(\eta)=d} \mathbb{I}\Big(t(\eta) \le t < t(\eta) + l^0(\eta) - l^-(\eta)\Big)\Big),$$

where $l^0(\eta)$ denotes the initial duration of η and $l^-(\eta)$ is the total reduction in the duration of η that is performed in lines 8 and 14 of procedure FDS. For each $\eta \in \mathcal{C}$, we have $0 \le l^-(\eta) \le l^0(\eta)$. Then,

$$\begin{aligned} v(\mathcal{C}) = \sum_{m=1}^{\mathcal{M}} \sum_{d=1}^{\mathcal{D}_m} \sum_{t=1}^{T} f\Big(r_{mdt} &- \sum_{\eta \in \mathcal{C}:\ m(\eta)=m,\ d(\eta)=d} \mathbb{I}\Big(t(\eta) \le t < t(\eta) + l^0(\eta)\Big) \\ &+ \sum_{\eta \in \mathcal{C}:\ m(\eta)=m,\ d(\eta)=d} \mathbb{I}\Big(t(\eta) + l^0(\eta) - l^-(\eta) \le t < t(\eta) + l^0(\eta)\Big)\Big). \end{aligned}$$

To simplify the notation, let

$$z_{mdt} := r_{mdt} - \sum_{\eta \in \mathcal{C}:\ m(\eta)=m,\ d(\eta)=d} \mathbb{I}\Big(t(\eta) \le t < t(\eta) + l^0(\eta)\Big),$$

$$\Delta z_{mdt} := \sum_{\eta\in\mathcal{C}:\ m(\eta)=m,\ d(\eta)=d} \mathbb{I}\Big(t(\eta)+l^0(\eta)-l^-(\eta)\le t < t(\eta)+l^0(\eta)\Big),$$

for all $m,\ d,\ t$. Therefore, we have

$$\begin{aligned} v(\mathcal{C}) &= \sum_{m=1}^{\mathcal{M}}\sum_{d=1}^{\mathcal{D}_m}\sum_{t=1}^{T} f(z_{mdt}+\Delta z_{mdt}) \\ &= \sum_{m=1}^{\mathcal{M}}\sum_{d=1}^{\mathcal{D}_m}\sum_{t=1}^{T}\left(f(z_{mdt})+\sum_{z'=z_{mdt}+1}^{z_{mdt}+\Delta z_{mdt}}\big(f(z')-f(z'-1)\big)\right) \\ &= v_{\text{ADDP}}+\sum_{m=1}^{\mathcal{M}}\sum_{d=1}^{\mathcal{D}_m}\sum_{t=1}^{T}\sum_{z'=z_{mdt}+1}^{z_{mdt}+\Delta z_{mdt}}\big(f(z')-f(z'-1)\big). \end{aligned}$$

Note, to obtain the second line, we add and subtract $f(z_{mdt}),\ f(z_{mdt}+1),\ldots,f(z_{mdt}+\Delta z_{mdt}-1)$. To obtain the third line, we apply the definition of v_{ADDP} given in equation (2.11).

Recall we assume r_{mdt} is sufficiently large such that $\sum_{g=1}^{G} x_{mdgt}\le r_{mdt}$, for all $m,\ d,\ t$; hence, $z_{mdt}\ge 0$, for all $m,\ d,\ t$. Moreover, $\Delta z_{mdt}\ge 0$, for all $m,\ d,\ t$, because the value of the function $\mathbb{I}(\cdot)$ is either 0 or 1. Finally, $z_{mdt}+\Delta z_{mdt}\le r_{mdt}$, for all $m,\ d,\ t$. Consider the definition of δ and note $\delta\ge f(z)-f(z-1)$, for all $1\le z\le r_{mdt}$, for all $m,\ d,\ t$. Then, we have

$$\begin{aligned} v(\mathcal{C}) - v_{\text{ADDP}} &\le \sum_{m=1}^{\mathcal{M}}\sum_{d=1}^{\mathcal{D}_m}\sum_{t=1}^{T}\sum_{z'=z_{mdt}+1}^{z_{mdt}+\Delta z_{mdt}}\delta \\ &= \delta\sum_{m=1}^{\mathcal{M}}\sum_{d=1}^{\mathcal{D}_m}\sum_{t=1}^{T}\left(\sum_{z'=z_{mdt}+1}^{z_{mdt}+\Delta z_{mdt}} 1\right) \\ &= \delta\sum_{m=1}^{\mathcal{M}}\sum_{d=1}^{\mathcal{D}_m}\sum_{t=1}^{T}\Delta z_{mdt} \\ &= \delta\sum_{m=1}^{\mathcal{M}}\sum_{d=1}^{\mathcal{D}_m}\sum_{t=1}^{T}\left(\sum_{\eta\in\mathcal{C}:\ m(\eta)=m,\ d(\eta)=d}\mathbb{I}\Big(t(\eta)+l^0(\eta)-l^-(\eta)\le t<t(\eta)+l^0(\eta)\Big)\right) \end{aligned}$$

$$
\begin{aligned}
&= \delta \sum_{\eta\in\mathcal{C}} \sum_{t=1}^{T} \mathbb{I}\Big(t(\eta) + l^0(\eta) - l^-(\eta) \le t < t(\eta) + l^0(\eta)\Big) \\
&= \delta \sum_{\eta\in\mathcal{C}} \left(\sum_{t=t(\eta)+l^0(\eta)-l^-(\eta)}^{t(\eta)+l^0(\eta)-1} 1 \right) \\
&= \delta \sum_{\eta\in\mathcal{C}} l^-(\eta).
\end{aligned}
$$

Define $l_1^-(\eta)$ and $l_2^-(\eta)$ as the reduction in the duration of η that is performed in lines 8 and 14 of procedure FDS, respectively. Then, $l^-(\eta) = l_1^-(\eta) + l_2^-(\eta)$. Therefore, we have:

$$
v(\mathcal{C}) - v_{\text{ADDP}} \;\le\; \delta \left(\sum_{\eta\in\mathcal{C}} l_1^-(\eta) + \sum_{\eta\in\mathcal{C}} l_2^-(\eta) \right). \tag{2.13}
$$

The following lemma is used in the remainder of the proof. Note that, for each $\eta \in \mathcal{C}$, the value of $g(\eta)$ is known and $g(\eta) \in \{1, \ldots, G\}$.

Lemma 2. *Let $\mathcal{C}^A$ be a feasible solution of ADDP and let $\mathcal{C}$ denote a solution generated by applying procedure FDS to $\mathcal{C}^A$. Then, $\mathcal{C}$ satisfies*

(a) $\sum_{\eta\in\mathcal{C}:\ g(\eta)=g',\ m(\eta)=m} l^0(\eta) - \sum_{\eta\in\mathcal{C}:\ g(\eta)=g'',\ m(\eta)=m} l^0(\eta) \le L$, *for all* g', g'', m, *and*

(b) $\sum_{\eta\in\mathcal{C}:\ g(\eta)=g,\ m(\eta)=m} l^0(\eta) \le \breve{h}_m + L$, *for all* g, m.

Lemma 2 can be proven similar to [FDA22], and hence is skipped. Part (b) of Lemma 2 implies:

$$
\begin{aligned}
& \sum_{\eta\in\mathcal{C}:\ g(\eta)=g,\ m(\eta)=m} l_1^-(\eta) \le L, \quad \forall g,\ m \\
\Rightarrow \quad & \sum_{m=1}^{\mathcal{M}} \sum_{g=1}^{G} \left(\sum_{\eta\in\mathcal{C}:\ g(\eta)=g,\ m(\eta)=m} l_1^-(\eta) \right) \le \sum_{m=1}^{\mathcal{M}} \sum_{g=1}^{G} L \\
\Rightarrow \quad & \sum_{\eta\in\mathcal{C}} l_1^-(\eta) \le \mathcal{M}GL.
\end{aligned} \tag{2.14}
$$

Moreover, we have:

$$\sum_{\eta\in\mathcal{C}:\ g(\eta)=g,\ m(\eta)=m} \left(l^0(\eta) - l_1^-(\eta)\right) \le \breve{h}_m, \quad \forall g,\ m \tag{2.15}$$

$$\Rightarrow \quad \sum_{m=1}^{\mathcal{M}} \left(\sum_{\eta\in\mathcal{C}:\ g(\eta)=g,\ m(\eta)=m} \left(l^0(\eta) - l_1^-(\eta)\right) \right) \le \sum_{m=1}^{\mathcal{M}} \breve{h}_m, \quad \forall g \tag{2.16}$$

$$\Rightarrow \quad \sum_{\eta\in\mathcal{C}:\ g(\eta)=g} \left(l^0(\eta) - l_1^-(\eta)\right) \le \sum_{m=1}^{\mathcal{M}} \breve{h}_m, \quad \forall g. \tag{2.17}$$

Because of line 14 of procedure FDS, we have:

$$\sum_{\eta\in\mathcal{C}:\ g(\eta)=g} l_2^-(\eta) = \left[\left(\sum_{\eta\in\mathcal{C}:\ g(\eta)=g} \left(l^0(\eta) - l_1^-(\eta)\right) \right) - H \right]^+, \quad \forall g \tag{2.18}$$

$$\Rightarrow \quad \sum_{\eta\in\mathcal{C}:\ g(\eta)=g} l_2^-(\eta) \le \left(\sum_{m=1}^{\mathcal{M}} \breve{h}_m \right) - H, \quad \forall g, \tag{2.19}$$

because of inequality (2.17) and noting that naturally we assume $H < \sum_{m=1}^{\mathcal{M}} \breve{h}_m$, because otherwise, the contractual limits on the total hours per group becomes redundant (which is not a well-designed DLCC contract).

Lemma 3. *$\sum_{\eta\in\mathcal{C}:\ g(\eta)=g} \left(l^0(\eta) - l_1^-(\eta)\right) \le H + \mathcal{M}L$, for all g.*

Proof of Lemma 3. Using Lemma 2(a), we have

$$\sum_{\eta\in\mathcal{C}:\ g(\eta)=g',\ m(\eta)=m} l^0(\eta) - \sum_{\eta\in\mathcal{C}:\ g(\eta)=g'',\ m(\eta)=m} l^0(\eta) \le L, \quad \forall g',\ g'',\ m$$

$$\Rightarrow \quad \sum_{m=1}^{\mathcal{M}} \left(\sum_{\eta\in\mathcal{C}:\ g(\eta)=g',\ m(\eta)=m} l^0(\eta) \right) - \sum_{m=1}^{\mathcal{M}} \left(\sum_{\eta\in\mathcal{C}:\ g(\eta)=g'',\ m(\eta)=m} l^0(\eta) \right) \le \mathcal{M}L, \quad \forall g',\ g''$$

$$\Rightarrow \quad \sum_{\eta\in\mathcal{C}:\ g(\eta)=g'} l^0(\eta) - \sum_{\eta\in\mathcal{C}:\ g(\eta)=g''} l^0(\eta) \le \mathcal{M}L, \quad \forall g',\ g''. \tag{2.20}$$

Next, we use a contradiction to prove Lemma 3. Assume g' exists such that

$$\begin{aligned} & \sum_{\eta\in\mathcal{C}:\ g(\eta)=g'} \left(l^0(\eta) - l_1^-(\eta)\right) > H + \mathcal{M}L \\ \Rightarrow \quad & \sum_{\eta\in\mathcal{C}:\ g(\eta)=g'} l^0(\eta) > H + \mathcal{M}L. \end{aligned} \tag{2.21}$$

Using Equations (2.20) and (2.21), we have

$$\begin{aligned} & \sum_{\eta\in\mathcal{C}:\ g(\eta)=g} l^0(\eta) > H, \quad \forall g \\ \Rightarrow \quad & \sum_{\eta\in\mathcal{C}} l^0(\eta) = \sum_{g=1}^{G} \left(\sum_{\eta\in\mathcal{C}:\ g(\eta)=g} l^0(\eta) \right) > GH, \end{aligned}$$

which contradicts the following constraint in ADDP:

$$\sum_{g=1}^{G}\sum_{m=1}^{\mathcal{M}}\sum_{d=1}^{\mathcal{D}_m}\sum_{t=1}^{T} x_{mdgt} \le GH.$$

Hence, the proof of Lemma 3 is complete. □

Using Lemma 3 and equation (2.18), we obtain $\sum_{\eta\in\mathcal{C}:\ g(\eta)=g} l_2^-(\eta) \le \mathcal{M}L$, for all g. Moreover, by applying equation (2.19), we find

$$\begin{aligned} & \sum_{\eta\in\mathcal{C}:\ g(\eta)=g} l_2^-(\eta) \le \min\left\{\mathcal{M}L,\ \left(\sum_{m=1}^{\mathcal{M}} \breve{h}_m\right) - H\right\}, \quad \forall g \\ \Rightarrow \quad & \sum_{\eta\in\mathcal{C}} l_2^-(\eta) = \sum_{g=1}^{G} \left(\sum_{\eta\in\mathcal{C}:\ g(\eta)=g} l_2^-(\eta) \right) \le G \min\left\{\mathcal{M}L,\ \left(\sum_{m=1}^{\mathcal{M}} \breve{h}_m\right) - H\right\}. \end{aligned} \tag{2.22}$$

Combining equations (2.13), (2.14), and (2.22), we obtain

$$v(\mathcal{C}) - v_{\text{ADDP}} \ \le \ \delta G \left(\mathcal{M}L + \min\left\{\mathcal{M}L,\ \left(\sum_{m=1}^{\mathcal{M}} \breve{h}_m\right) - H\right\}\right),$$

and hence, the proof of Theorem 2 is complete.

2.10.4 Proof of Theorem 3

Proof of part (a). KMP$^\circ$ minimizes $\sum_{m=1}^{\mathcal{M}} p_{m,3}$ subject to $\sum_{m=1}^{\mathcal{M}} p_{m,1} \leq GK$, $\sum_{m=1}^{\mathcal{M}} p_{m,2} \leq GH$, and $\mathbf{p}_m \in conv(\mathbb{P}_m)$, for all m. We use a contradiction. Let $\{\mathbf{p}_m^\circ\}_m$ denote an optimal solution, and assume m' exists such that $\mathbf{p}_{m'}^\circ$ is not on the lower envelop $\mathcal{E}_{m'}$. Thus, $\hat{\mathbf{p}}_{m'} \in conv(\mathbb{P}_m)$ exists such that $\hat{\mathbf{p}}_{m'} = (p_{m',1}^\circ, p_{m',2}^\circ, p_{m',3}^\circ - a)$, for some $a > 0$. In the optimal solution, if we replace $\mathbf{p}_{m'}^\circ$ with $\hat{\mathbf{p}}_{m'}$, the new solution is feasible and its objective value is $\sum_{m=1}^{\mathcal{M}} p_{m,3}^\circ - a$, which is strictly smaller than the optimal value.

Proof of part (b). Let $\{\nu_{m,i}^\circ\}_{m,i}$ denote an optimal solution for KMP$^\circ$. We prove an integral optimal solution for KMP$^\circ$ exists. Define $p_{m,\ell}^\circ := \sum_{i=1}^{N_m} p_{m,i,\ell} \nu_{m,i}^\circ$, for $\ell = 1, 2$, which denote the number of calls and hours, respectively, assigned to month m. Define $\mathcal{S}_1^\circ := GK - \sum_{m=1}^{\mathcal{M}} p_{m,1}^\circ$ and $\mathcal{S}_2^\circ := GH - \sum_{m=1}^{\mathcal{M}} p_{m,2}^\circ$ as the slack of constraints (2.4b) and (2.4c), respectively. If $\mathcal{S}_1^\circ$ is fractional, some fractional $p_{m,1}^\circ$ exists. By increasing such fractional $p_{m,1}^\circ$'s and decreasing $\mathcal{S}_1^\circ$, we obtain an alternative optimal solution with integral $\mathcal{S}_1^\circ$ (because of Theorem 3(a) and that $\mathcal{E}_m$ is non-increasing in the number of calls). Repeating this procedure for $\mathcal{S}_2^\circ$, we obtain an optimal solution with integral $\mathcal{S}_1^\circ$ and $\mathcal{S}_2^\circ$. Consider the following problem:

$$\min \quad \sum_{m=1}^{\mathcal{M}} \mathcal{E}_m(u_{m,1}, u_{m,2}) \tag{2.23}$$

$$\text{s.t.} \quad \sum_{m=1}^{\mathcal{M}} u_{m,\ell} = \sum_{m=1}^{\mathcal{M}} p_{m,\ell}^\circ, \quad \forall \ell = 1, 2, \tag{2.24}$$

$$\lfloor p_{m,\ell}^\circ \rfloor \leq u_{m,\ell} \leq \lceil p_{m,\ell}^\circ \rceil, \quad \forall m, \ \ell = 1, 2, \tag{2.25}$$

where $u_{m,\ell}$ is the decision variable. By adding slacks to constraints (2.25) and performing some algebraic manipulations, constraints (2.25) is equivalent to constraints (2.26)-(2.28).

$$u_{m,\ell} + \bar{\mathcal{S}}_{m,\ell,1} = \lceil p_{m,\ell}^\circ \rceil, \quad \forall m, \ \ell = 1, 2, \tag{2.26}$$

$$\bar{\mathcal{S}}_{m,\ell,1} + \bar{\mathcal{S}}_{m,\ell,2} = \lceil p^{\circ}_{m,\ell} \rceil - \lfloor p^{\circ}_{m,\ell} \rfloor, \quad \forall m, \ \ell = 1, 2, \tag{2.27}$$

$$\bar{\mathcal{S}}_{m,\ell,1}, \ \bar{\mathcal{S}}_{m,\ell,2} \geq 0, \quad \forall m, \ \ell = 1, 2. \tag{2.28}$$

Therefore, we minimize objective function (2.23) over the feasible set defined by constraints (2.24) and (2.26)-(2.28). The coefficient matrix of constraints (2.24), (2.26), and (2.27) is *totally unimodular*, and the right-hand-side values of these constraints are integral. Thus, all extreme points are integral. Because, for all m, the lower envelop $\mathcal{E}_m$ is linear over the unit square $[k, k+1] \times [h, h+1]$, for all $k \in \{0, 1, \ldots, G\breve{k}_m - 1\}$ and $h \in \{0, 1, \ldots, G\breve{h}_m - 1\}$, the objective function is linear over the feasible set (because constraints (2.25) bounds the feasible set by a unit square for each month). Thus, an extreme point is optimal. Therefore, we proved an optimal solution exists with integral numbers of calls and hours for each month.

Proof of part (c). Note we assume, for each month, all combinations of calls and hours have been enumerated; hence, for each m, $\mathbf{p}_{m,i_m}$ exists such that $p_{m,i_m,1}$ and $p_{m,i_m,2}$ are equal to the (integral) optimal numbers of calls and hours, respectively. We construct an integral optimal solution for KMP$^{\circ}$ by setting $\nu_{m,i_m} = 1$, for all m, and $\nu_{m,i} = 0$, for all m and $i \neq i_m$ (because we assume all $\mathbf{p}_{m,i}$'s are on the lower envelops). Thus, we proved an integral optimal solution exists for KMP$^{\circ}$. This optimal solution is feasible for KMP and provides the same objective value. Therefore, recalling KMP$^{\circ}$ is a relaxation of KMP, the integral optimal solution that we obtained for KMP$^{\circ}$ is also optimal for KMP, and hence, the proof is complete.

2.10.5 Proof of Proposition 2

We prove this proposition using a reduction from the *3-partition problem* [GJ79], similar to [FDA22]. Consider the 3-partition problem with the set of $3n$ elements $A = \{a_1, \ldots, a_{3n}\}$, such that $a_i \in \mathbb{Z}^+$, $\frac{B}{4} < a_i < \frac{B}{2}$, for all i, and $\sum_{i=1}^{3n} a_i = nB$. Consider an instance of WP where $G = n$, and each group has three calls and B hours. Let the ECPs in the month consist of $3n$ days with profiles that exactly match the elements of the 3-partition problem. The

3-partition problem has a feasible solution if A can be partitioned into G subsets of size B, in which case the corresponding ECPs can be assigned to the G groups, where each group receives three calls and B hours.

2.10.6 Proof of Theorem 4

DA allocates calls on a rotational basis to the groups and a call is randomly selected. A call assigned to a group is executed only if the group has the necessary hours. To derive an upper bound on the expected hours lost by DA, we ignore the AMP (KMP) limit on hours and assume each group will be called $\frac{k_m^*}{G}$ times. The loss in optimality depends on the number of hours that are wasted because some calls were not executed due to lack of hours. If a group reaches its hour limit on the τ^{th} call, the hours assigned on subsequent $\frac{k_m^*}{G} - \tau$ calls are wasted, and a portion of the τ^{th} call may also be wasted. Note τ is the hitting time for the sum of hours assigned to a group to reach or exceed the available hours $\frac{h_m^*}{G}$, which is defined as

$$\tau := \min \left\{ t \in \{0, 1, \dots\} \;\middle|\; \sum_{i=1}^{t} l_i \geq \frac{h_m^*}{G} \right\}.$$

The expected number of calls wasted is

$$\mathbf{E}\left[\frac{k_m^*}{G} - \tau \;\middle|\; \tau \leq \frac{k_m^*}{G} \right].$$

Therefore, by Wald's identity, the expected hours lost is

$$\bar{l}\left(\mathbf{E}\left[\frac{k_m^*}{G} - \tau \;\middle|\; \tau \leq \frac{k_m^*}{G} \right] + \epsilon \right),$$

where $\epsilon \in [0,1)$ accounts for the shortages on the τ^{th} call. Because the total hours available per group are $\frac{k_m^*}{G}\bar{l}$, the expected percentage lost is

$$\frac{G}{k_m^*}\left(\mathbf{E}\left[\frac{k_m^*}{G}-\tau \;\middle|\; \tau \leq \frac{k_m^*}{G}\right]+\epsilon\right) = \frac{1}{n}\left(\mathbf{E}\left[n-\tau \mid \tau \leq n\right]+\epsilon\right),$$

where we let $n = \frac{k_m^*}{G}$. By [Sie68, Theorem 1], we have

$$\lim_{n\to\infty} \frac{\tau - n}{\sqrt{n}\sigma_l(\bar{l})^{-1}} \xrightarrow{\mathscr{D}} \mathcal{N}(0,1).$$

The expected number of calls unused is equivalent to the number of units salvaged in a news-vendor problem when the stocking level is equal to the mean. Hence, in the limit we get

$$\lim_{n\to\infty} \mathbf{E}(n-\tau|\tau \leq n) \;\to\; \beta\sqrt{n}\sigma_l\frac{1}{\bar{l}}.$$

The proof is completed by applying Wald's identity to derive the expected hours lost and dividing by the hours available to get the following limit on the percentage of hours lost as the number of calls (n) increases:

$$\lim_{n\to\infty} \frac{\beta\sigma_l}{\sqrt{n}\,\bar{l}} + \frac{\epsilon}{n},$$

and hence, DA is asymptotically optimal.

2.11 Appendix C: Numerical Experiments and Figures

In this section, we provide some additional figures that we've omitted from the numerical experiments section due to space considerations. In particular, Figure 2.21 compares the number of consumed calls and hours by the AMP-HWP and KMP-HWP approaches. Moreover, Figure 2.22 plots the difference between the number of consumed calls and hours by these two approaches across different months. On the other hand, Figure 2.23 illustrates the

number of allocated and consumed calls by AMP-HWP and KMP-HWP to different months. Finally, Figure 2.24 presents the same plot for the number of hours. As the figures indicate, the variation of the allocated resources is generally higher for AMP-HWP in comparison with KMP-HWP. Furthermore, KMP-HWP keeps most of the resources for peak months, while AMP starts to use the resources earlier. We observe that the allocation patterns of the resources are also different, as discussed in the body of the chapter.

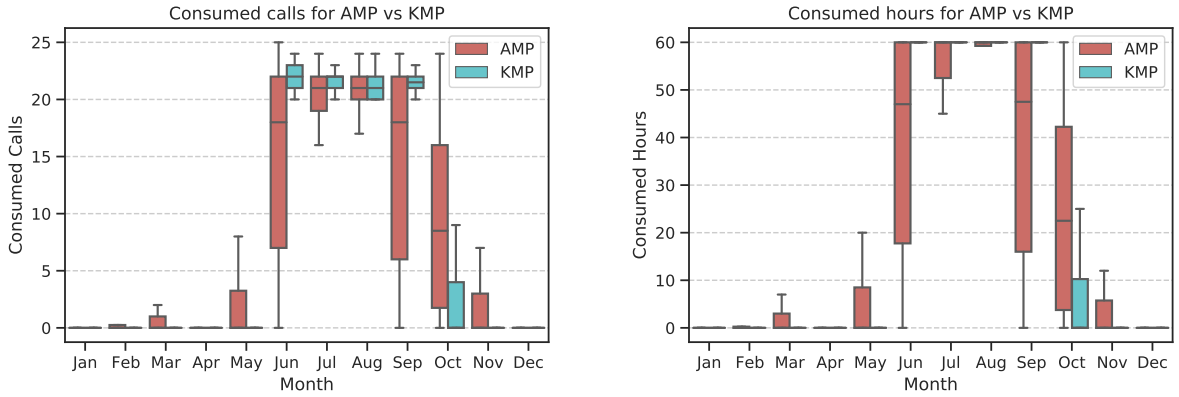

(a) Consumed calls (b) Consumed hours

Figure 2.21: Comparison of the number of consumed calls and hours by AMP and KMP in different months.

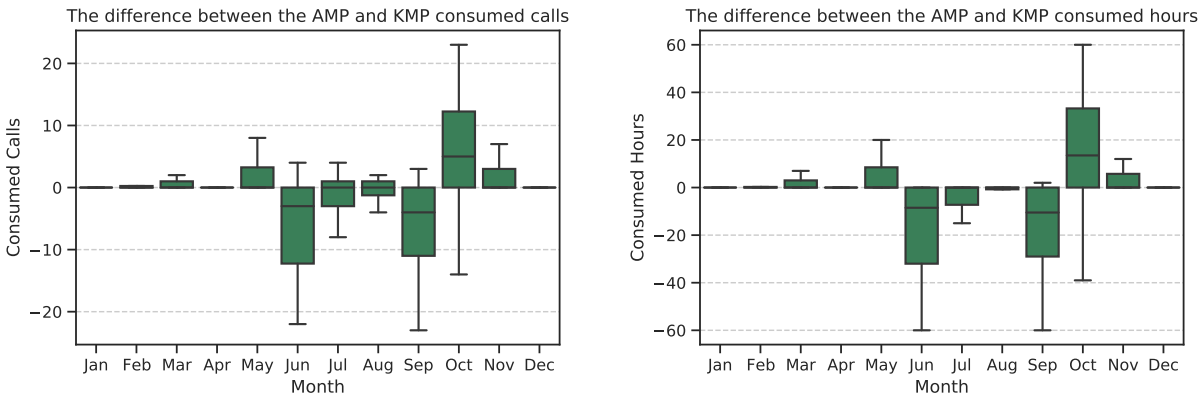

(a) Consumed calls (b) Consumed hours

Figure 2.22: The difference between the number of consumed calls and hours by AMP and KMP in different months.

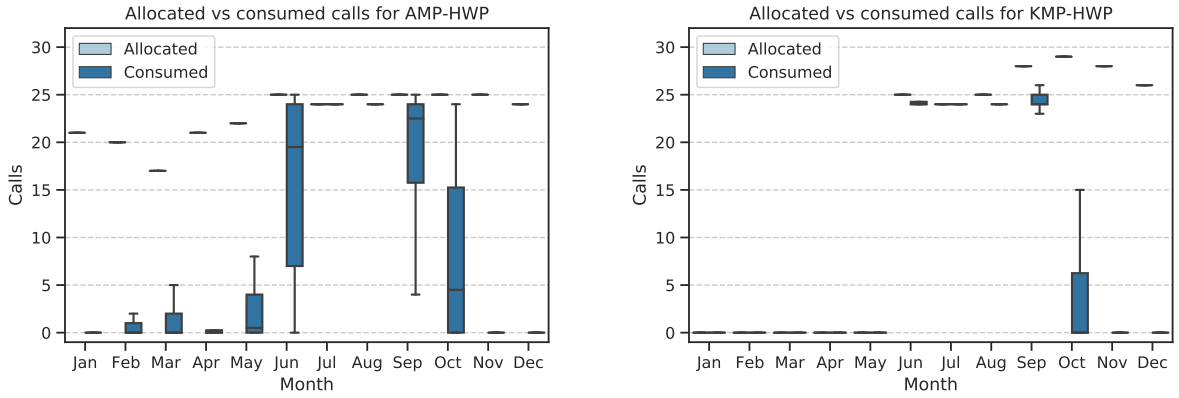

(a) AMP-HWP (b) KMP-HWP

Figure 2.23: The number of allocated and consumed calls for AMP-HWP and KMP-HWP.

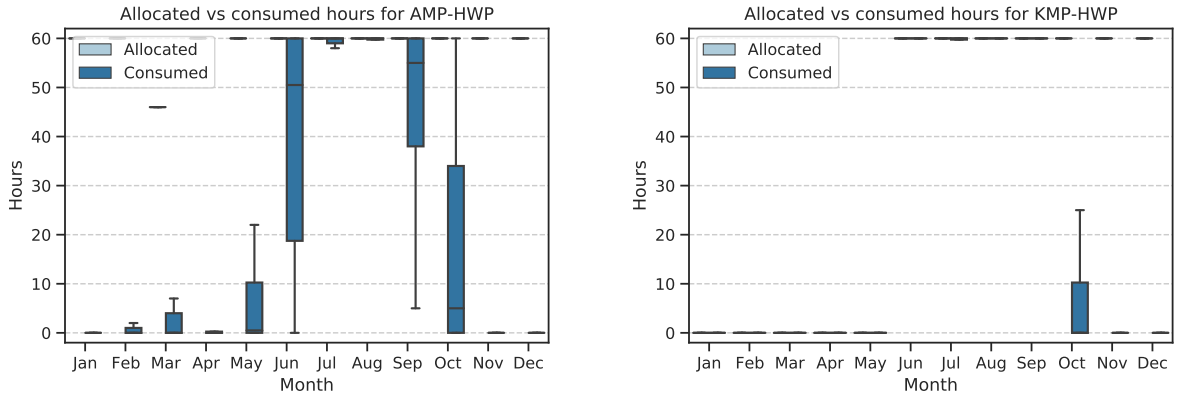

(a) AMP-HWP (b) KMP-HWP

Figure 2.24: The number of allocated and consumed hours for AMP-HWP and KMP-HWP.

CHAPTER 3

A Framework for Modeling the Choice of Healthcare Intervention Strategies

3.1 Introduction

As the public safety-net healthcare service provider for all the Riverside County's residents, the Riverside University Health System (RUHS) aims to improve the population-level quality of care as well as access to health care services. The RUHS's behavioral health department, which offers community-based services to patients with behavioral health problems (i.e. mental health conditions or substance abuse disorders), annually reaches to more than 50,000 people across the Riverside County. Having a professional team of approximately 1,000 full-time employees including clinicians, psychiatrists, peer specialists, and paraprofessionals, RUHS offers services in three broad categories of mental health, substance use, and public guardian programs primarily to individuals on Medi-Cal and Medicare insurance plans.

Their adult services are divided into eleven mental health clinic sites, six of which accommodate older adults, along with three urgent care clinics and two inpatient facilities. The adult clinics provide services including but not limited to psychiatric assessment, crisis intervention, recovery management, case management, and medication services. Moreover, housing assistance, outreach to mentally-ill homeless individuals, and inpatient hospitalization are also part of their services for adult patients. On the other hand, the children's mental health clinics work closely with schools to provide services to children who encounter academic problems as a result of severe emotional issues. The programs serving older adults focus

on wellness, recovery, and resiliency of the customers in order to help them recover and maintain a physical and emotionally healthy lifestyle. Preventive mental health services are also provided to the at-risk population for homelessness, institutionalization, and substance abuse.

Our research has been conducted in collaboration with administrative members of RUHS and especially aimed at exploring opportunities for sustainable improvement of the population's health in the long-run. In the following, we briefly discuss the key ideas that motivated our research, introduce the core components of our model, justify various aspects of our modeling approach based on the characteristics of the case under study, and finally take a look at the problem from a wider perspective. We also explain the mathematical structure of our decision-making problem. Moreover, we will see that to estimate the parameters of the optimization problem, we need to design a statistical disease progression model as well.

The main idea behind the current work was initially inspired by observing the key role that decisions regarding the trade-off between two contradictory approaches for providing healthcare services to a population play in the long-term health and well-being of the population. One approach is providing care to patients as needed and another approach is providing early intervention services. As [PPH16] explain, almost half of all the healthcare expenditures in the US are associated with only the top 5% of the healthcare users. Consequently, implementing appropriate preventive intervention strategies can potentially reduce the burden on society significantly. On the other hand, the median age-of-onset for many life-time mental disorders is known to be around the early 20s [KAA07]. Moreover, there is often a long delay of almost one decade, on average, between the onset of symptoms and the beginning of intervention [KCD05]. Therefore, preventive interventions play a crucial role, especially in the behavioral healthcare sector. The impact of preventive interventions on the progression of behavioral health disorders has been well studied in the literature [ADM18]. However, as mentioned, there is an inherent trade-off in making decisions regarding the choice of intervention strategies. Although improving the general health of the population by implementing early intervention

strategies can reduce the corresponding potential expenditure in the future, it incurs an additional cost as well which may not be desirable in resource-constrained systems.

Our modeling mechanism is supposed to capture the impact of decisions regarding the choice of intervention on the dynamics of the population's health and provide a quantitative framework for evaluating long-term costs and benefits of such decisions. However, we do not restrict the set of interventions to only preventive and regular interventions since such distinction may not always be possible in practice. From a high-level perspective, we categorize individuals into multiple groups and then model the impact of intervention strategies on their disease progression paths. Our disease progression model doesn't assume direct access to the underlying health state of patients. Therefore, we use the physician observations for categorizing the patient. We will briefly introduce a common approach for classifying patients based on their functioning-level in the next section. The model is supposed to determine the optimum assignment of available interventions to patients. For example, an optimal solution to our problem may include assigning 40% of patients who are in the so-called moderate condition at a given time period to an intervention package comprising case management and medication as well as assigning the remaining 60% to another intervention package consisting from regular psychotherapy and partial hospitalization. The impacts of these intervention options on the disease progression of patients are then evaluated by estimating the probability according to which patients transition into other health states after receiving each intervention package. More specifically, a transition probability matrix can be estimated for any arbitrary intervention package which indicates the proportion of the population that transition from each category to each category during one period of time as a result of receiving this intervention. We prefer a population-based modeling approach that studies the aggregate impact of high-level decisions on the whole population. This way of looking at the system is in contrast to modeling patient-level impacts of interventions and then adding up the effects. We will later discuss the rationale behind this choice in the next sections.

To estimate the population-level transition rates of the planning problem, we need to have

a model for the patient-level progression of disease. For this purpose, we employ a Continuous-Time Hidden Markov Model (CT-HMM) which consists of an underlying Continuous-Time Markov Chain (CTMC) for modeling the true health state of patient as well as an emission model for generating the observational data based on the patient's health state. We've chosen a latent variable model since direct information on the patients' underlying health state is rarely available in the behavioral healthcare domain. Therefore, physician observation data is used instead as a proxy for making inferences on the true health state. Moreover, the continuous-time nature of this model allows us to describe the progression of disease as a continuous variable using irregular discrete-time observations and nicely handles the missing data problem.

We define a linear objective function for the planning problem. To understand the motivation behind this choice, here we present a brief discussion on the revenue sources of the system under study. The RUHS's state funding is based on Medi-Cal 1991 and 2011 realignments that cover involuntary care and state hospitals. The Mental Health Services Act (MHSA) also provides funding sources for outpatient care, full-service partnership programs, and capital investment for programs covering the incarcerated, unemployed, and homeless population. However, the administrators report an estimated annual $150 million funding shortfall. Certain categories of services, such as those provided by Federally Qualified Health Centers (FQHCs), are cost-reimbursed and represent a volume-dependent revenue opportunity.

In our framework, the healthcare service provider is assumed to receive a certain amount of budget and is supposed to consume that budget on keeping a population of patients as healthy as possible. The patients may require additional healthcare services, depending on their condition, and incur some cost to the system. We divide the services that patients may need into two major categories: 1- non-hospitalization services, and 2- hospitalization service. Such a distinction between these two types of services has its roots in the payment system under which RUHS operates. Specifically, our model assumes that payment to the provider is performed based on a hybrid Fee For Service (FFS)-capitation system. In this payment

scheme, non-hospitalization services are reimbursed according to an FFS payment system while a capitation payment system is used for hospitalization services. The provider tries to maximize the net payoff associated with the entire population of patients. We'll later discuss the link between the hybrid FFS-capitation payment system and the linearity of the objective function. We use a cost measure as the objective function of our problem rather than the common quality of life measures, such as Quality-Adjusted Life Years (QALY), due to some inherent weaknesses of such measures which will be discussed in detail later.

Our model also incorporates resource consumption constraints. The mood disorders, schizophrenia, and mental disorders related to substance abuse are the most prevalent diseases and conditions among RUHS's patients. These patients require a wide array of services such as general psychiatry, trauma psychiatry, therapy for anxiety and mood disorders, addiction medication, case management, group therapy, and neuropsychological testing. Each service requires certain resources ranging from human resources (e.g. psychiatrists, social workers, and nursing staff) to physical resources (e.g. hospital beds). We emphasize that the nature and severity of patient's condition significantly affect the type and frequency of service needs. As an example, general psychotherapy sessions often work well for early-stage depression, while patients suffering from acute schizophrenia or mood disorders (such as major depression and bipolar disorder) might even need help around the clock. We model the resource requirement and budget constraints by appropriately defining resource and cost matrices. More specifically, for any given resource we consider a matrix that indicates the amount of this resource required for providing a certain intervention to a patient in a certain category. Such a matrix is then used for defining linear inequality constraint that provides an upper bound on the total amount of this resource being consumed by any arbitrary solution to the optimization problem. For example, an intervention package including partial or full hospitalization may require one unit of hospital beds for patients in the so-called sever category and half a unit of hospital beds for patients in the so-called moderate category during one period of time. Similarly, a cost matrix is defined for the budget constraint. We

do not impose any methodological restrictions on the type of resource that a service may need in our model, however, we assume that the resource requirement parameters are known and homogeneous among patients in a certain class of health state.

The inefficiency issues go beyond a certain sector of the healthcare industry and are not limited to the geographical area of our study. It is well-known that although the US per-capita spending on healthcare is more than any other developed country, this spending has not resulted in better health outcomes such as life expectancy and the prevalence of chronic conditions [GS08]. [GS18] conduct a critical review study on the evidence of the inefficiency of the US healthcare system. The authors explore the fundamental differences between the healthcare industry and competitive industries in which inefficient players will naturally be eliminated from the market and put emphasis on the significance of reforming all the organizational, financial, and operational aspects of the healthcare system as the only way of improving the system-wide efficiency.

Although there is no theoretical limitation in our framework which prevents its application to other healthcare sectors, the body of our research makes the most sense in the context of modeling behavioral healthcare systems. In fact, the chronic nature of most of the behavioral health conditions makes long-term planning for the choice of intervention strategies more meaningful. The problem is still worth studying since behavioral health disorders are common among the US populations. The reports provided by the National Institute of Mental Health (NIMH) indicate that more than 46 million US adults experienced mental illness in the year 2017 and almost 11 million of them suffered from a Serious Mental Illness (SMI). The anxiety disorder, clinical depression, Attention Deficit Hyperactivity Disorder (ADHD), bipolar disorder, Post-Traumatic Stress Disorder (PTSD), schizophrenia, and autism are the most common mental disorders in the United States. Mental health disorders can potentially have adverse impacts on all the social, physical and economic aspects of the patient's life and generally result in a lower quality of life (see [DW11], [BB11], and [LDD17]). According to a report by the Substance Abuse and Mental Health Services Administration (SAMHSA), the

total U.S. mental health treatment spending is expected to reach approximately $240 billion in 2021 which is 5.5% of expected nationwide health spending [SAM14]. We refer the readers to the study by [Kna03] for more details about the hidden costs of mental health disorders for the economy.

In general, the U.S. behavioral health care system is under-resourced. More specifically, psychiatry is one of the top workforce shortages among all medical specialties and the number per capita has a decreasing trend. [SNS18] estimate between 14,280 to 31,091 psychiatrist workforce shortage in the year 2024. The authors further explain that although a slow expansion is expected to happen after 2025, it is not clear whether it can resolve the shortage in the next 30 years. Moreover, as a consequence of deinstitutionalization policies in the U.S. mental health care system [SJ13], the number of public psychiatric beds fell from 340 to 17 beds per 100,000 population during the period 1955 to 2005 [TEG15]. A recent report by the National Survey on Drug Use and Health (NSDUH) indicates that only 42% of adults with a mental disorder received mental health services in the year 2017 [SAM18]. This, in turn, highlights the importance of system-level efficiency improvements which make serving more patients in the current under-resourced situation possible.

As an extension of the baseline disease progression model, we also present a more sophisticated version of our CT-HMM model that allows for occasional observation of the underlying health state. We will discuss the motivations behind presenting this extension by highlighting the differences in the nature of diseases such as cancer and heart disease with behavioral health conditions. Finally, we present a patient-level intervention planning problem based on the extended model, which employs a single patient's historical data for decision-making.

The main contributions of our work are as follows. First, we develop a CT-HMM disease progression model that incorporates the effect of healthcare interventions, and present an efficient approach for learning the model parameters. We employ this model for estimating the transition rates between different levels, which is a basis for population-level planning.

Second, we present an extension of the model with occasionally-observed underlying health state. We then discuss how to estimate the underlying health state for a given patient based on their historical data. Such estimation would be helpful for individual-level intervention planning.

3.2 Literature Review

In this section, we briefly review the relevant literature in three major areas, each concerning one aspect of our work: overview of the US healthcare system, disease progression modeling, and healthcare intervention planning. In the first part, we start by providing a high-level background on the traditional frameworks for economic evaluation of healthcare interventions. Afterward, we mention some of the largest and most successful public health programs that aim to improve the efficiency of healthcare systems in the US and across the globe. Finally, we go over the fundamentals of healthcare payment systems to better understand the motivations behind our choice of the planning objective function. In the disease progression part, we mention some of the key characteristics of healthcare conditions that make development of statistical models challenging. We also explain the motivation behind choosing CT-HMM models for our work, and highlight some complex aspects of disease progression that our model doesn't cover. Lastly, we discuss the healthcare intervention planning literature at both the population-level and patient-level. We refer the readers to [BHP09] for a review of the literature on modeling and simulation in healthcare.

3.2.1 An Overview of the US Healthcare System

Economic analysis of healthcare interventions. A branch of the health economics literature focuses on evaluating the costs and outcomes associated with healthcare interventions and making comparisons among alternative treatment options. At first, notice that the distinction between effectiveness and cost-effectiveness of healthcare interventions should be recognized

[RBK05]. An intervention is said to be effective if it helps to improve clinical outcomes such as disease symptoms and health-related quality of life measures. However, the goal of economic analysis is to determine whether an intervention is worth its costs.

[Kna19] provides an extensive literature review on the economic analysis of behavioral health interventions. The author conceptualizes three levels of decision-making questions: 1- "how to best treat a specific disorder?", 2- "how to allocate the available resources within the health care sector?", and 3- "how to allocate the available resources across different areas of public policy?". The way cost is measured is almost the same among all the levels, however, outcome measurement requires totally different approaches. A certain class of economic analysis tools is suggested by the author for each level of decision making questions. The Cost-Effectiveness Analysis (CEA) is the most appropriate approach for making decisions regarding specific health needs (e.g. choosing the best treatment option for a patient with depression). Limiting the scope of the problem to one disease or condition makes it possible to use patient-centered measures of clinical outcomes for the analysis [GCW11]. For example, the Incremental Cost-Effectiveness Ratio (ICER) statistic for two given treatment options is defined as the ratio of the difference in costs to the difference in health outcomes. This ratio indicates how much cost the system incurs for each unit of additional benefit that the patient receives if a certain alternative treatment option is provided. At a higher level, the Cost-Utility Analysis (CUA) methods are suggested for modeling planning and resource allocation in the public health sector. A well-defined measure of utility is though needed to make comparisons across outcomes associated with different diseases and conditions possible. Finally, the Cost-Benefit Analysis (CBA) methods which are supposed to inform the highest level of decisions, involve comparisons among the different areas of public policy. Therefore, an economic analysis at this level requires a monetary representation for clinical outcomes which may not be always available [Fre10]. We refer the readers to [MVP11] for a critical review of recent studies on economic analysis of preventive interventions for mental health disorders.

Among the aforementioned three layers of decision making, the first and second are closer to our problem under study. However, we look at the problem from a broader perspective, in the sense that we set our ultimate goal as minimizing the total cost incurred to the system in the long-run by the entire population. Specifically, we've decided to use a cost measure instead of health-related quality of life measures, such as QALY, which have been commonly used as generic utility measures for cost-utility analyses by economics and healthcare researchers [Tor06]. In the following, we discuss the rationale behind this choice. Despite all the advantages that QALY measures have, such as comparability and explicitness, many studies cast doubt on whether QALY measures are appropriate for healthcare planning and resource allocation purposes. [CHK97] critically explore various aspects of QALY measures and highlight the inherent weaknesses of this class of measures that can potentially make them inappropriate for planning purposes. Specifically, they raise ethical questions such as "who (e.g. the general public, clinicians, politicians, etc.) is responsible for making value judgments concerning the dimensions of well-being to be covered and valuation of health states?". Moreover, the authors bring out an underlying assumption required for using QALYs for planning purposes which may not be widely agreed upon: "a healthy year of life expectancy is taken to be of equal value to everyone irrespective of their characteristics". On the other hand, there are some methodological challenges for using QALYs at the planning level as well. For example, the QALY scores, measured under certain circumstances, are not easily generalizable to other localities, different case-mixes, different medical practices, etc. Last but not least, the authors explain that mental health care interventions are aimed more at improving life than extending or saving a life. Therefore, providing behavioral interventions generally produces less QALYs than interventions that significantly extend both life expectancy and quality of life which means that maximizing health gains in this way is likely to result in the marginalization of such interventions.

The public health perspective. To make the discussion more concrete, let's review the problem from the perspective of healthcare professionals and public health administrators. The

World Health Organization (WHO) recognizes performing economic analysis as an essential step within the process of identifying an efficient allocation of resources in healthcare systems. However, they emphasize that such analysis is insufficient and needs to be systematically weighed up against other objectives (e.g. fairness considerations) as well as implementation issues (e.g. feasibility, acceptability, and sustainability) [Org06]. [MCP13] critically review the challenges which are known to prevent practical usage of cost-effectiveness analysis for resource allocation and priority setting in healthcare. Specifically, the authors refer to the lack of generalizability, limited policy relevance, methodological confounding, and implementation issues as main factors that have contributed to the limited practical impact of cost-effectiveness studies on decisions making procedures. They also conduct case studies on six real-world large-scale examples of priority setting programs in the context of behavioral healthcare management, including the WHO's Choosing Interventions that are Cost-Effective (CHOICE) project, Oregon's state Medicaid program called Oregon Health Plan (OHP), and the Mental Health Economics European Network (MHEEN) program. The authors conclude that adopting a narrow economic perspective is a common element among those projects that failed to achieve their goal. Moreover, they emphasize the key role that offering a strong technical foundation and including considerations outside the scope of economic analysis play in the success of such programs. [EC08] review the literature on the use of economic evaluation in local decision-making (i.e. both the meso and micro levels of management). The authors add to the aforementioned list some other factors that have prevented the practical usage of economic analysis, specifically in the US healthcare system. For example, they refer to some institutional and political factors such as inflexibility of healthcare budgets, as well as some cultural factors such as a general tendency for prioritizing efficiency over cost-effectiveness. Indeed, the successful implementation of any planning framework, including ours, highly depends on incorporating these considerations in the model.

Healthcare payment models. To understand the cost-based tradeoffs that healthcare providers such as RUHS face, we need to be familiar with their funding mechanisms. The

FFS payment system, under which the provider gets reimbursed proportional to the service cost, is still one of the most common healthcare payment models in the US. It is well known that the FFS payment system incentivizes a high volume of care rather than encouraging providers to improve efficiency. In fact, many healthcare professionals and researchers believe that the FFS payment system is the biggest obstacle in improving healthcare delivery [PK16]. Capitation payment is one of the most well-known alternatives for the FFS system. Under this payment scheme, the insurer pays a fixed amount, often to a single organization, for providing services to a patient during an entire period of time. This payment system has also been widely criticized, especially for bearing the wrong risk to the providers. In the bundled payment system, which is often referred to as a key component of recent healthcare reforms, the provider is paid for managing a certain condition or disease over a pre-specified period of time. The rationale behind this idea is that such a payment model gives provider incentives for improving the quality of care since no additional payment will be performed for any kind of complications that may happen during the aforementioned period of time. As explained beforehand, RUHS operates under a payment system in which certain categories of services are reimbursed according to the FFS payment scheme while a capitation payment system is used for the rest of the service types.

[AMN17] compares the performance of a traditional FFS payment system with bundled payment in a simple setting with one provider who makes decisions regarding only the intensity of treatment. The authors highlight the financial risk that the provider bears as the main weakness of this model and propose a hybrid FFS-bundled payment scheme as a better alternative. [GTW19] compare the FFS and bundled payment systems in terms of their impact of social welfare in a public healthcare system using a game-theoretic model. [STY19] adopt a similar modeling approach for analyzing a hybrid payment system under Yardstick competition. In modeling healthcare systems, it is often assumed that there are three parties including an insurer, a provider, and a set of beneficiaries. In the literature, a three-stage Stackelberg model is often used for capturing the impacts of interactions among

these parties. However, here we do not study the game-theoretic aspect of the problem. We will later elaborate on the mathematical details of the payment system under study.

3.2.2 Disease Progression Modeling

In this section, we mention the unique challenges of developing statistical models for disease progression and then briefly review a few common State Transition Models (STMs). In the next section, we will explain how these models have been used by the operations management community for planning and resource allocation purposes. We refer the readers to [PFU18] for a review of patient flow modeling approaches that we do not cover, such as queueing models which are commonly used in analyzing hospital operations.

In the past couple of decades, standard practices for recording Electronic Health Records (EHR) have significantly improved, and more data with a higher accuracy is now available as a consequence. Developing mathematical models for disease progression has been a topic of interest for researchers in a variety of fields, including operations management and statistics. Our focus in this research project is on incorporating the effect of healthcare interventions on disease progression. Such a statistical model would then allow for formulating an optimization framework for choosing the best set of interventions under certain constraints at both the population and individual levels. Developing accurate models for disease progression is inherently a challenging task. The authors explain some aspects of this complexity in [WSW14] very well:

- **Unobserved health condition**: What physicians observe during regular visits is usually a noisy signal of the underlying disease state. In other words, accessing the true health state of the patients requires accurate examination, which is often risky, expensive, or even impossible (in certain domains such as behavioral healthcare).

- **Discrete and irregularly-spaced observation times**: Indeed, the underlying disease progression happens continuously over time. However, physicians usually can only

access discrete-time snapshots of this process. In many real-world EHR datasets, even the duration of the time between consecutive visits varies, which makes traditional discrete-time models inappropriate for capturing the trends. Therefore, a continuous-time disease progression model is needed that is trained based on irregularly-spaced discrete-time observations.

- **The effect of interventions and covariates**: Classical disease progression models often assume that there are no external factors that can affect the course of disease progression. However, in many situations the intervention decisions by the physician or covariates such as age and gender can have an impact on how the disease progresses, and the model should be able to incorporate these factors.

Biology papers in areas such as pharmacology often develop linear regression models for tracking the changes of certain disease indicators over time or as a function of certain factors [CB16, Mou12]. In the literature, long-term modeling of diseases based on EHR data is often performed using state transition models. These models either assume the patients share the same characteristics or control for these factors. A key issue in working with state transition models is that sometimes it is difficult to represent the disease state using a limited number of discrete states [CBS12].

There is an extensive body of literature on using Discrete-Time Markov Chains (DTMCs) for modeling behavioral conditions (e.g. depression [AKG12]). On the other hand, CTMCs that consider a continuous-time underlying process and add more complexity to the model structure in this way have also been used in the past [GLL94]. There are many variations of these models that each of them tries to improve upon the baseline models by relaxing a certain assumption or adding some level of flexibility [LM13]. Hidden Markov Models (HMMs) are an extended class of Markovian models in which the underlying process is assumed to be unobserved. Hence, only a noisy signal is observed that is related to the underlying variable via a statistical distribution. Specifically, CT-HMMs are constructed

based on continuous-time underlying variables that are observed at discrete points in time, which provides a great degree of flexibility for describing real-world data [LLL15].

Different aspects of these models have also been extended in the past [WSW14]. For example, the Markovian assumption that the transitions of the underlying disease state depend only on the current state of the system might be restrictive in some applications in which the history (e.g. the duration of stay) also plays a role in the disease dynamics, and some alternative approaches have been proposed for overcoming this limitation [SCS14, AV18, AS19]. Moreover, [AHS17] considers a model in which the physician decisions affect the duration of time between consecutive visits. Training a CT-HMM disease progression model using discrete-time irregular observations has been done in the past [LM13]. In this work, the impact of covariates has also been incorporated into the model using the proportional hazards approach. However, the main drawback of their parameter estimation method is that the generator matrix is assumed to be diagonalizable in all the EM steps. We refer the readers to [SGL19] for a similar paper that employs CT-HMMs for modeling healthcare data. Most of the disease progression models that have been proposed in the literature ignore the effect of healthcare interventions on the disease dynamics. However, there are some papers that incorporate the impact of healthcare interventions on the physician observation. For example, [SCS20] designs a DT-HMM with continuous observation values in which interventions shift the distribution of the observations. The main advantage of this methodological framework is that it allows for handling heterogeneous patient-specific intervention effects.

An underlying assumption in our model is that patients can be appropriately divided into distinct groups. The concept of classification of patients based on their service needs has been traditionally used in the healthcare industry for planning and reimbursement purposes. For example, the Diagnosis-Related Groups (DRG) is widely used in hospital payment systems in all the states and other countries as well. However, there is a body of literature that argues diagnosis-based measures, such as DRG, are poor indicators of resource utilization by patients, if considered alone [McC95]. In the context of behavioral healthcare,

measures for the functioning level of patients (e.g. community living skills, negative social behavior, affective distress, etc.) are often referred to as the best tools for determining the service needs and appropriate levels of care. As an example, [USS03] propose an empirical approach for clustering patients into a few groups of similar individuals based on their characteristics, clinical history, and functioning-level. The authors then define levels of care (e.g. brief intervention, medication, rehabilitation, etc.) and manually assign each class to an appropriate level of care. We refer the readers to [BHC16] for an exhaustive review of the functioning-level assessment instruments. The operations management community has also developed resource allocation and planning models based on the notion of levels of care. A classic example of such efforts is the work by [LDG86] in which the authors define a linear programming model for assigning predefined service packages to patient groups. They used a set of basic functioning level measures for clustering patients into various classes and manually determine the set of allowable service packages for each class. In this research, we do not restrict ourselves to a certain categorization mechanism and keep the model at the conceptual level.

3.2.3 Healthcare Intervention Planning

In the following, we present a brief review of the intervention planning literature. At first, we discuss some of the traditional approaches that have been proposed for population-level planning. Afterward, we go over some patient-level planning models. In particular, we introduce Markov Decision Processes (MDPs) as the most common framework for formulating sequential decision-making problems. We then discuss how CT-HMMs can be used as a baseline for estimating the initial state of POMDPs.

Population-level planning. The operations management community has previously used state transition models for characterizing progression of disease at the population-level. In this approach, a representative patient goes through a process with Markovian dynamics and the probabilistic quantities of the resulting model will be used as estimators of the actual

quantities. For example, [DIJ13] uses this approach for providing capacity allocation guidelines in community-based healthcare delivery for chronic diseases. A stochastic optimization problem then maximizes a measure of population health (e.g. QALYs) in a multi-period decision-making setting. A similar approach has been proposed for modeling the optimal screening problem for Hepatocellular Carcinoma by [LLV19]. Furthermore, Markovian models have also been used in the literature for characterizing the visit patterns of patients, which is crucial part in intervention policy outcome assessment and modeling [KVM11].

Despite all the advantages of population-level modeling framework in operations management applications such as [CBD07], we have not found extensive use cases for these approaches in the healthcare decision-making research literature [ES12]. As an example of the population-based approaches, [DRR15] studies the problem of HIV screening and intervention. Specifically, their method divides patients into different groups based on their health state and assumes the transition rates between these groups (in a multi-period setting) is constant. Following a more complex methodological framework, [HLZ19] formulates an approximate dynamic programming problem and propose a multi-fidelity rollout algorithm for solving the multi-period policy optimization problem. Other disease progression modeling techniques such as using differential equations have also been employed in the literature for describing the population-level spread of diseases such as HIV [ALB14]. In this paper, an optimization problem then helps to make the best resource allocation decisions in order to keep the population as healthy as possible based on a certain metric. Similarly, [NS17] studies the problem of allocating inventory in humanitarian health settings using a population-based model.

Patient-level planning. The traditional approach for incorporating the effect of healthcare decisions and interventions on the course of disease is to formulate the problem using MDPs rather than HMMs. Therefore, here we briefly review the reinforcement learning literature and discuss the advantages and limitations of this approach. In the past few years, intervention planning using RL methods has gained the attention of the healthcare research community

[FMD20, GJM18, LSS20]. One of the most common frameworks for modeling the healthcare sequential decision-making process is Dynamic Treatment Regimes (DTR) [CM14]. For more details regarding the DTR approach as well as other applications of RL in healthcare modeling and planning, please refer to [SLL11] and [YLN21]. Among the RL methods, the framework developed by [FHD20] is the most similar one to our work. A discrete-time partially-observed MDP with continuous observation values is used here for modeling the progression of Sepsis disease. The authors suggest an EM algorithm for learning the parameters of their model similar to an Input-Output HMM [BF94]. As the reinforcement learning framework was originally developed for online settings that allow for active exploration, training models in an offline setting using fixed datasets is a tricky task and requires special attention. We refer the readers to [LKT20] for an overview of the offline reinforcement learning literature. Specifically, safety is a critical issue in the healthcare domain in offline settings and there is a body of literature on modifying RL techniques for such areas in which making mistakes can be extremely risky and costly [FKS21].

The main advantage of employing POMDPs is that they naturally incorporate the effect of healthcare decisions into the disease dynamics. On the other hand, the main disadvantage of the RL framework is that most of the traditional RL algorithms are specialized for discrete-time decision making applications. Although we face a discrete-time decision making problem in our healthcare problems, we are also interested in capturing the true dynamics of the underlying disease process, which is continuous-time, using irregularly-spaced discrete-time observations. In the literature, there exist some complex methods for overcoming these limitations of POMDPs [ASK20]. However, in this research work we prefer to use CT-HMMs for modeling the healthcare data as the capabilities of these models for capturing the complexities of EHR data have already been proven.

3.3 Efficient Parameter Learning for CT-HMM Disease Progression Model

As part of formulating the intervention planning problem, we need to estimate the transition rates of the patients between different health states. We will explain later in details how a disease progression model can be used for this purpose. In the following, we setup the foundation for learning CT-HMMs using discrete-time irregularly observed data based on a technique that was developed by [LLL15].

3.3.1 Model Setup

Consider a dataset that includes N trajectories, each corresponding to the physician observations of a certain disease. Denote by I and J, respectively, the number of discrete values that the underlying health state and the physician observation can take. Moreover, let T_n be the number of observations in the trajectory of patient n and $\tau_{n1}, \cdots, \tau_{nT_n}$ be the observation times. Define $\mathbf{y}_n^{\tau_{ns}} \in \{0, \cdots, J-1\}$ (for $s = 1, \cdots, T_n$) and $\mathbf{z}_n^{\tau} \in \{1, \cdots, I\}$ (for the time index $\tau \in \mathbb{R}_+$) as the random variables that correspond to the observation and the underlying latent state for patient n. Notice that $\mathbf{z}$ is continuous time, while $\mathbf{y}$ is discrete-time, as explained earlier. In addition, define $\tilde{\mathbf{T}}_n$ as the actual number of transitions of the underlying health state, i.e. $\mathbf{z}_n^{\tau}$. Similarly, suppose $\tilde{\boldsymbol{\tau}}_{ns}$ (for $1 \leq s \leq \tilde{\mathbf{T}}_n$) indicate the corresponding transition times.

Let $\mathcal{Q}$ and $\mathcal{P}(\Delta\tau)$ be $I \times I$ matrices that correspond to the generator matrix and the transition probability function of the CT-HMM, respectively. The transition probability functions for time period $\Delta\tau$ will be $\mathcal{P}(\Delta\tau) = \exp(\Delta\tau\mathcal{Q}) = \sum_{l=0}^{\infty} (\Delta\tau)^l \frac{\mathcal{Q}^l}{l!}$. Here the model is assumed to be time-homogeneous and power series expansion is used for defining the matrix exponential. Furthermore, we define $\mathcal{E}(i,j) = \mathbb{P}\{\mathbf{y}_n^{\tau_{ns}} = j | \mathbf{z}_n^{\tau_{ns}} = i\}$ (for any $s = 1, \cdots, T_n$) as the observation model. Assuming a binomial emission distribution with probability μ_i for state i, we have $\mathcal{E}(i,j) = \binom{J-1}{j} \mu_i^j (1-\mu_i)^{J-1-j}$ for $1 \leq i \leq I, 0 \leq j \leq J-1$.

3.3.2 Expectation Maximization Algorithm for Parameter Learning

Parameter learning for latent variable models is often performed using the EM algorithm [Jac11]. [MHS07] presents an efficient technique for learning the parameters of CTMCs based on discrete-time and irregularly-spaced observations. In this section, we'll review a learning algorithm for CT-HMMs based on this idea that was proposed by [LLL15].

Suppose we have a homogeneous dataset in terms of the parameters $\mathcal{Q}$ and $\mathcal{E}$. The complete-data likelihood for patient n can be represented in terms of $\mathbf{y}_n^{\tau_{n\,1:T_n}}$ $(1 \leq s \leq \tilde{T}_n)$ and $\mathbf{z}_n^{\tau}$ in the following way:

$$
\begin{aligned}
&\mathfrak{L}_n(\mathcal{Q}, \mathcal{E} | \mathbf{y}_n^{\tau_{n\,1:T_n}}, \mathbf{z}_n^{\tau}) \triangleq \mathbb{P}\{\mathbf{y}_n^{\tau_{n\,1:T_n}}, \mathbf{z}_n^{\tau} | \mathcal{Q}, \mathcal{E}\} \\
&= \prod_{\tilde{s}=1}^{\tilde{T}_n - 1} \left[\frac{\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\,\tilde{s}+1}}]}{-\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\tilde{s}}}]} \times (-\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\tilde{s}}}])\, e^{\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\tilde{s}}}]\,(\tilde{\tau}_{n\,\tilde{s}+1} - \tilde{\tau}_{n\,\tilde{s}})} \right] \times \prod_{s=1}^{T_n} \mathcal{E}(z_n^{\tau_{ns}}, y_n^{\tau_{ns}}) \\
&= \prod_{\tilde{s}=1}^{\tilde{T}_n - 1} \left[\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\,\tilde{s}+1}}] \times e^{\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\tilde{s}}}]\,(\tilde{\tau}_{n\,\tilde{s}+1} - \tilde{\tau}_{n\,\tilde{s}})} \right] \times \prod_{s=1}^{T_n} \mathcal{E}(z_n^{\tau_{ns}}, y_n^{\tau_{ns}})
\end{aligned}
$$

where the transition probability of the underlying DTMC from $z_n^{\tilde{\tau}_{n\tilde{s}}}$ to $z_n^{\tilde{\tau}_{n\,\tilde{s}+1}}$ is reflected in $\frac{\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\,\tilde{s}+1}}]}{-\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\tilde{s}}}]}$, and $-\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\tilde{s}}}]\, e^{\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\tilde{s}}}]\,(\tilde{\tau}_{n\,\tilde{s}+1} - \tilde{\tau}_{n\,\tilde{s}})}$ is the sojourn time probability distribution of state $z_n^{\tilde{\tau}_{n\tilde{s}}}$. Denote by $\boldsymbol{\chi}_{n\,ik}$ $(1 \leq i, k \leq I, k \neq i)$ and $\boldsymbol{\psi}_{n\,i}$ $(1 \leq i \leq I)$ the total number of transitions and the total amount of time spent in each state by the underlying health state variable. We can simplify the complete-data likelihood as:

$$
\mathfrak{L}_n(\mathcal{Q}, \mathcal{E} | \mathbf{y}_n^{\tau_{n\,1:T_n}}, \mathbf{z}_n^{\tau}) = \prod_{i=1}^{I} \prod_{k=1, k\neq i}^{I} \mathcal{Q}[i,k]^{\chi_{n\,ik}} \times \prod_{i=1}^{I} e^{\mathcal{Q}[i,i]\,\psi_{n\,i}} \times \prod_{s=1}^{T_n} \mathcal{E}(z_n^{\tau_{ns}}, y_n^{\tau_{ns}})
$$

For the dataset of N patients, the total complete-data likelihood is the multiplication of these terms, i.e. $\mathfrak{L} = \prod_{n=1}^{N} \mathfrak{L}_n$. Clearly, we can not evaluate this function due to the existence of latent variable $\mathbf{z}$. Therefore, we assume an initial value $(\tilde{\mathcal{Q}}, \tilde{\mathcal{E}})$ for the set of parameters is given and evaluate the expected log-likelihood $\mathbb{E}_{\tilde{\mathcal{Q}}, \tilde{\mathcal{E}}}[\log \mathfrak{L}(\mathcal{Q}, \mathcal{E} | \mathbf{y}_n^{\tau_{n\,1:T_n}})]$. Afterward, we

maximize this function and update the parameters. The EM iterations will continue until convergence. We can expand the expected log-likelihood function as:

$$
\begin{aligned}
\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}[\log \mathfrak{L}(\mathcal{Q},\mathcal{E}|\mathbf{y}_n^{\tau_n\,1:T_n}] &= \mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\left[\sum_{n=1}^{N}\sum_{i=1}^{I}\sum_{k=1,k\neq i}^{I} \boldsymbol{\chi}_{n\,ik}\,\log(\mathcal{Q}[i,k])\,\middle|\,\mathbf{y}_n^{\tau_n\,1:T_n}\right] \\
&+ \mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\left[\sum_{n=1}^{N}\sum_{i=1}^{I}\mathcal{Q}[i,i]\,\boldsymbol{\psi}_{n\,i}\,\middle|\,\mathbf{y}_n^{\tau_n\,1:T_n}\right] + \mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\left[\sum_{n=1}^{N}\sum_{s=1}^{T_n}\log(\mathcal{E}(\mathbf{z}_n^{\tau_{ns}},\mathbf{y}^{\tau_{ns}}))\,\middle|\,\mathbf{y}_n^{\tau_n\,1:T_n}\right] \\
&= \sum_{n=1}^{N}\sum_{i=1}^{I}\sum_{k=1,k\neq i}^{I}\left(\log(\mathcal{Q}[i,k])\times\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\left[\boldsymbol{\chi}_{n\,ik}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\right]\right) \\
&+ \sum_{n=1}^{N}\sum_{i=1}^{I}\left(\mathcal{Q}[i,i]\times\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\left[\boldsymbol{\psi}_{n\,i}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\right]\right) + \sum_{n=1}^{N}\sum_{s=1}^{T_n}\left(\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\left[\log(\mathcal{E}(\mathbf{z}_n^{\tau_{ns}},y_n^{\tau_{ns}}))\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\right]\right)
\end{aligned}
$$

By maximizing this function with respect to $\mathcal{Q}$ and $\mathcal{E}$, we can update the parameters. In particular, the solution of the generator matrix is:

$$
\hat{\mathcal{Q}}[i,k] = \frac{\sum_{n=1}^{N}\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\left[\boldsymbol{\chi}_{n\,ik}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\right]}{\sum_{n=1}^{N}\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\left[\boldsymbol{\psi}_{n\,i}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\right]};\ (k\neq i) \qquad \hat{\mathcal{Q}}[i,i] = -\sum_{k=1,k\neq i}^{I}\hat{\mathcal{Q}}[i,k]
$$

The main challenge would then be to calculate the expected values of the $\boldsymbol{\chi}$ and $\boldsymbol{\psi}$ summary statistics. One of the most common approaches for calculating these quantities for simpler HMM models is to use the Monte Carlo Expectation-Maximization (MCEM) method in which a set of samples of the latent variable are generated based on the current parameter values as well as the observed data, and then the expectations are approximated by their samples averages [LC01]. In section 3.11.1 in the appendix, we consider a DT-HMM model and describe a Gibbs sampling algorithm for generating samples from the posterior of the graph. Afterward, we discuss how the expectations of the summary statistics can be approximated using the generated samples. However, sample generation for CT-HMM models with a complex structure is computationally challenging. In the literature, some approaches for calculating these expectations have been proposed for situations in which the generator matrix is diagonalizable [MHS07, WSW14], which is usually not the case in EM iterations.

In the following, we'll go over a recently proposed method that overcomes this issue.

3.3.3 An Efficient Approach for Evaluating the Expectations

In [LLL15], the authors develop an efficient technique for calculating these expectations. At first, they define $\boldsymbol{\chi}_{n\,ik}^{\tau_{n\,(s)}}$ and $\boldsymbol{\psi}_{n\,i}^{\tau_{n\,(s)}}$ as the corresponding values for the time period $\tau_{n\,(s)} = (\tau_{ns}, \tau_{n\,s+1}]$ and condition on the latent variables:

$$\begin{aligned}
\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\chi}_{n\,ik}\,|\,\mathbf{y}_n^{\tau_{n\,1:T_n}}\big] &= \sum_{s=1}^{T_n-1}\sum_{\tilde{i}=1}^{I}\sum_{\tilde{k}=1}^{I}\mathbb{P}\{\mathbf{z}_n^{\tau_{n\,s}}=\tilde{i},\mathbf{z}_n^{\tau_{n\,s+1}}=\tilde{k}\,|\,\mathbf{y}_n^{\tau_{n\,1:T_n}}\}\\
&\qquad\times\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\chi}_{n\,ik}^{\tau_{n\,(s)}}\,|\mathbf{z}_n^{\tau_{n\,s}}=\tilde{i},\,\mathbf{z}_n^{\tau_{n\,s+1}}=\tilde{k}\big]\\
\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\psi}_{n\,i}\,|\,\mathbf{y}_n^{\tau_{n\,1:T_n}}\big] &= \sum_{s=1}^{T_n-1}\sum_{\tilde{i}=1}^{I}\sum_{\tilde{k}=1}^{I}\mathbb{P}\{\mathbf{z}_n^{\tau_{n\,s}}=\tilde{i},\mathbf{z}_n^{\tau_{n\,s+1}}=\tilde{k}\,|\,\mathbf{y}_n^{\tau_{n\,1:T_n}}\}\\
&\qquad\times\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\psi}_{n\,i}^{\tau_{n\,(s)}}\,|\mathbf{z}_n^{\tau_{n\,s}}=\tilde{i},\,\mathbf{z}_n^{\tau_{n\,s+1}}=\tilde{k}\big]
\end{aligned}$$

Hence, the E-step reduces to efficient calculation of the end-state conditioned expectations $\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\chi}_{n\,ik}^{\tau_{n\,(s)}}\,|\mathbf{z}_n^{\tau_{n\,s}}=\tilde{i},\,\mathbf{z}_n^{\tau_{n\,s+1}}=\tilde{k}\big]$, and $\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\psi}_{n\,i}^{\tau_{n\,(s)}}\,|\mathbf{z}_n^{\tau_{n\,s}}=\tilde{i},\,\mathbf{z}_n^{\tau_{n\,s+1}}=\tilde{k}\big]$ as well as the posterior probability $\mathbb{P}\{\mathbf{z}_n^{\tau_{n\,s}}=\tilde{i},\mathbf{z}_n^{\tau_{n\,s+1}}=\tilde{k}\,|\,\mathbf{y}_n^{\tau_{n\,1:T_n}}\}$. The posterior probability can be calculated by constructing a time-inhomogeneous DT-HMM from the CT-HMM. Specifically, the transition probability matrix of the DT-HMM between two given visits $[\tau_s, \tau_{s+1})$ will be the transition probability function of the CT-HMM, i.e. $\mathcal{P}(\tau_{s+1}-\tau_s) = \exp((\tau_{s+1}-\tau_s)\tilde{\mathcal{Q}})$. Therefore, evaluating the posterior probability can be performed by calculating the probability of the state sequence of the DT-HMM.

The remaining tricky part is the calculation of the expectations. To do so, the authors suggest using a technique for representing the quantities as integrals based on the work proposed by [HJ05] and [HJ11]:

$$\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\chi}_{n\,ik}^{\tau_{n\,(s)}}\,|\mathbf{z}^{\tau_s}=\tilde{i},\,\mathbf{z}^{\tau_{s+1}}=\tilde{k}\big] = \frac{\tilde{\mathcal{Q}}_{ik}}{(e^{(\tau_{s+1}-\tau_s)\tilde{\mathcal{Q}}})_{\tilde{i}\tilde{k}}}\int_0^{(\tau_{s+1}-\tau_s)}(e^{x\tilde{\mathcal{Q}}})_{\tilde{i}i}(e^{(\tau_{s+1}-\tau_s-x)\tilde{\mathcal{Q}}})_{k\tilde{k}}\,dx \quad (3.1)$$

$$\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\left[\boldsymbol{\psi}_{n\,i}^{\tau_{n\,(s)}} \,|\mathbf{z}^{\tau_s}=\tilde{i},\,\mathbf{z}^{\tau_{s+1}}=\tilde{k}\right]=\frac{1}{(e^{(\tau_{s+1}-\tau_s)\tilde{\mathcal{Q}}})_{\tilde{i}\tilde{k}}}\int_0^{(\tau_{s+1}-\tau_s)}(e^{x\tilde{\mathcal{Q}}})_{\tilde{i}i}(e^{(\tau_{s+1}-\tau_s-x)\tilde{\mathcal{Q}}})_{i\tilde{k}}\,dx \quad (3.2)$$

Now, they use a classical trick for evaluating this convolution-like integral form. Specifically, suppose $\mathcal{B}$ is an $I \times I$ matrix and we're interested in evaluating $\int_0^t e^{x\tilde{\mathcal{Q}}}\,\mathcal{B}\,e^{(t-x)\tilde{\mathcal{Q}}}\,dx$. We can construct matrix $\mathcal{A}=\begin{pmatrix}\tilde{\mathcal{Q}} & \mathcal{B}\\ 0 & \tilde{\mathcal{Q}}\end{pmatrix}$ and calculate $e^{t\mathcal{A}}$. The result of the integral would be equal to the upper right corner of this matrix [Van78]. In our case, we choose $\mathcal{B}$ in a way that it has zeros in all the elements and one in element (i,k) and use the $(\tilde{i},\tilde{k})$'th element as the result of the integral.

3.4 Modeling Intervention

In this section, we reformulate the model to incorporating the effect of interventions and covariates on disease progression. Although the modeling framework is general, our motivating examples are behavioral conditions such as schizophrenia and bipolar disorder. We know that for most behavioral conditions no information on the underlying health state is available to physicians. Therefore, we assume in our model that access to the corresponding variable is not possible and only a noisy signal in the form of physician observation data is available. These observations can be the result of questionnaires or other clinical evaluation methods. As discussed earlier, in these cases there are usually a set of intervention packages (e.g. medication, case management, hospitalization, etc.) that each might be more appropriate for certain categories of patients. Hence, the main trade-off is between the effectiveness of each intervention and the resource consumption associated with it. Notice that our model doesn't impose monotonicity assumption on the strength or aggressiveness of the interventions.

From the modeling perspective, we show how parameterizing the generator matrix of the CT-HMM using the proportional hazards method can be used for incorporating the impact of interventions on disease progression (Figure 3.1). In the following, we'll describe an EM

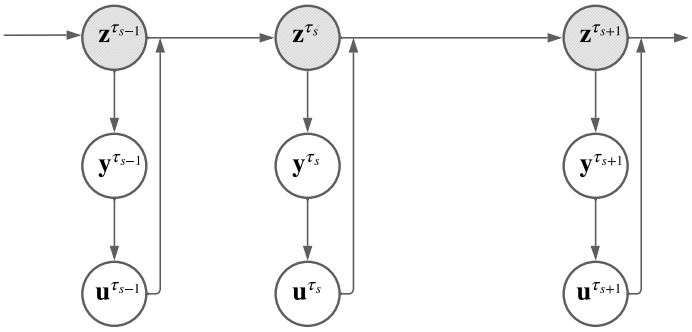

Figure 3.1: Incorporating the effect of interventions into the model by parameterizing the generator matrix.

algorithm for estimating the model parameters by employing the aforementioned ideas. We refer the readers to a recent work by the author of this dissertation for more details [Gho22].

3.4.1 Model Setup

Let $\mathbf{z}_n^{\tau}$ indicate the continuous-time underlying health state. We use the notations of $\mathbf{z}_n^{\tau_{ns}}$ and $\mathbf{z}_n^{\tau_{n\,(s)}}$ for the value of this variable at time point τ_{ns} and the entire interval $(\tau_{ns}, \tau_{n\,s+1}]$. Moreover, suppose $\mathbf{y}_n^{\tau_{ns}}$ is the physician observation variable at visit time τ_{ns}. An intervention option (including "no-intervention") is chosen based on the observation at each visit time and use $\mathbf{u}_n^{\tau_{ns}}$ for denoting the corresponding variable. We assume access to a dataset of N patient trajectories for $(\mathbf{y},\ \mathbf{u})$, each of length T_n. All the three variables $\mathbf{z}$, $\mathbf{y}$, and $\mathbf{u}$ take discrete values and we use the ranges $\{1, \cdots, I\}$, $\{0, \cdots, J-1\}$, and $\{0, \cdots, L-1\}$ for them. The main parameters of the model include the generator matrix $\mathcal{Q} \in \mathbb{R}^{I \times I}$, the emission distribution $\mathcal{E}(i,j) = \mathbb{P}(\mathbf{y}_n^{\tau_{ns}} = j \,|\, \mathbf{z}_n^{\tau_{ns}} = i)$, the intervention distribution $\mathcal{G}(j,l) = \mathbb{P}(\mathbf{u}_n^{\tau_{ns}} = l \,|\, \mathbf{y}_n^{\tau_{ns}} = j)$, and the initial health state distribution $\pi \in \mathbb{R}^I$. Both the emission and intervention distributions are assumed to be binomial with parameters μ_i $(1 \leq i \leq I)$ and η_j $(0 \leq j \leq J-1)$, respectively.

3.4.2 Parameterization of the Generator Matrix

In the context of behavioral healthcare, there are many covariates, such as age or gender, that influence the disease progression dynamics. Moreover, we're interested in incorporating the impact of healthcare interventions into the model, as discussed earlier. One approach for making the generator matrix of the CT-HMM a function of the interventions and covariates is the proportional hazards method [Cox72]. In particular, the elements of the generator matrix are defined to be exponential functions with the inputs being linear functions of the interventions and covariates [KL85]. In mathematical terms, let $w_n \in \mathbb{R}^D$ be the covariate vector. The generator matrix elements corresponding to each intervention l can be represented as:

$$[\mathcal{Q}]_{ikl} = \delta_{ikl} \, \exp(\rho_{ikl}^{'} w_n); \quad 1 \leq i, k \leq I, \, k \neq i, \, 0 \leq l \leq L-1 \tag{3.3}$$

$$[\mathcal{Q}]_{iil} = - \sum_{k=1, k \neq i}^{I} [\mathcal{Q}]_{ikl}; \quad 1 \leq i \leq I, \, 0 \leq l \leq L-1 \tag{3.4}$$

In the following, we use the notation $\mathcal{Q}_{n\,(s)}$ for the generator matrix in the time period $(\tau_s, \tau_{s+1}]$, which has $\delta_{ik\,\mathbf{u}_n^{\tau_{ns}}} \, \exp(\rho_{ik\,\mathbf{u}_n^{\tau_{ns}}}^{'} w_n)$ as its (i,k) elements.

3.4.3 EM Algorithm for Parameter Learning

Denote the set of parameters that we're interested to estimate by $\Theta \triangleq (\pi, \rho, \delta, \mu, \eta)$. Similar to the previous model, we can formulate the complete-data likelihood as:

$$\begin{aligned}\mathcal{L}_n(\Theta) = \mathbb{P}(\mathbf{u}_n^{\tau_{n\,1:T_n}}, \mathbf{y}_n^{\tau_{n\,1:T_n}}, \mathbf{z}_n^{\tau_n} | \Theta) = \mathbb{P}(\mathbf{u}_n^{\tau_{nT_n}} | \mathbf{y}_n^{\tau_{nT_n}}, \eta) \, \mathbb{P}(\mathbf{y}_n^{\tau_{nT_n}} | \mathbf{z}_n^{\tau_{nT_n}}, \mu) \\ \times \prod_{s=1}^{T_n - 1} \left[\mathbb{P}(\mathbf{z}_n^{\tau_{n\,(s)}} | \mathbf{z}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}}, \rho, \delta) \, \mathbb{P}(\mathbf{u}_n^{\tau_{ns}} | \mathbf{y}_n^{\tau_{ns}}, \eta) \, \mathbb{P}(\mathbf{y}_n^{\tau_{ns}} | \mathbf{z}_n^{\tau_{ns}}, \mu) \right] \times \mathbb{P}(\mathbf{z}_n^{\tau_{n1}})\end{aligned}$$

By taking the log of this function we get:

$$\log\, \mathcal{L}_n(\Theta) = \log \mathbb{P}(\mathbf{z}_n^{\tau_{n1}}) + \sum_{s=1}^{T_n-1} \left[\log \mathbb{P}(\mathbf{z}_n^{\tau_n\,(s)} | \mathbf{z}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}}, \rho, \delta) \right] + \sum_{s=1}^{T_n} \left[\log \mathbb{P}(\mathbf{u}_n^{\tau_{ns}} | \mathbf{y}_n^{\tau_{ns}}, \eta) + \log \mathbb{P}(\mathbf{y}_n^{\tau_{ns}} | \mathbf{z}_n^{\tau_{ns}}, \mu) \right] \tag{3.5}$$

Clearly, the complete-data log likelihood function includes the latent variable $\mathbf{z}$ and can not be evaluated. Hence, we assume $\tilde{\Theta} \triangleq (\tilde{\pi}, \tilde{\rho}, \tilde{\delta}, \tilde{\mu}, \tilde{\eta})$ is given as an initial guess, and follow an iterative approach in which the expectation of the log-likelihood function is calculated and then the parameters are updated by maximizing $\mathbb{E}[\log\, \mathcal{L}_n(\Theta) | \mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}]$.

Let $\chi_{n\,ik}^{\tau_n\,(s)}$ indicate the total number of transitions that the underlying health state variables makes between visits $(\tau_{ns}, \tau_{n\,s+1}]$. Similarly, denote by $\psi_{n\,i}^{\tau_n\,(s)}$ the sojourn time corresponding to the same interval. The transition part of the likelihood function can then be simplified as:

$$\mathbb{P}(\mathbf{z}_n^{\tau_n\,(s)} | \mathbf{z}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}}, \rho, \delta) = \prod_{i=1}^{I} \prod_{\substack{k=1 \\ k \neq i}}^{I} \left[Q_{n\,(s)} \right]_{ik}^{\chi_{n\,ik}^{\tau_n\,(s)}} \times \prod_{i=1}^{I} \exp\left([\mathcal{Q}_{n\,(s)}]_{ii}\, \psi_{n\,i}^{\tau_n\,(s)} \right)$$

$$\log \mathbb{P}(\mathbf{z}_n^{\tau_n\,(s)} | \mathbf{z}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}}, \rho, \delta) = \sum_{i=1}^{I} \sum_{\substack{k=1 \\ k \neq i}}^{I} \left(\chi_{n\,ik}^{\tau_n\,(s)} \log([\mathcal{Q}_{n\,(s)}]_{ik}) - \psi_{n\,i}^{\tau_n\,(s)}\, [\mathcal{Q}_{n\,(s)}]_{ik} \right)$$

For the emission and intervention parts of the function we have:

$$\log \mathbb{P}(\mathbf{y}_n^{\tau_{ns}} | \mathbf{z}_n^{\tau_{ns}}, \mu) = \sum_{i=1}^{I} \left[\log \binom{J-1}{\mathbf{y}_n^{\tau_{ns}}} + \mathbf{y}_n^{\tau_{ns}} \log \mu_i + (J - 1 - \mathbf{y}_n^{\tau_{ns}}) \log(1 - \mu_i) \right] \mathbb{I}\{\mathbf{z}_n^{\tau_{ns}} = i\}$$

$$\log \mathbb{P}(\mathbf{u}_n^{\tau_{ns}} | \mathbf{y}_n^{\tau_{ns}}, \eta) = \sum_{j=0}^{J-1} \left[\log \binom{L-1}{\mathbf{u}_n^{\tau_{ns}}} + \mathbf{u}_n^{\tau_{ns}} \log \eta_j + (L - 1 - \mathbf{u}_n^{\tau_{ns}}) \log(1 - \eta_j) \right] \mathbb{I}\{\mathbf{y}_n^{\tau_{ns}} = j\}$$

The initial state probability can also be written as $\log \mathbb{P}(\mathbf{z}_n^{\tau_{n1}}) = \sum_{i=1}^{I} \log \pi_i\, \mathbb{I}\{\mathbf{z}_n^{\tau_{n1}} = i\}$.

Hence, the total function will be:

$$
\begin{aligned}
\log \mathcal{L}(\Theta) = \sum_{n=1}^{N} \log\, \mathcal{L}_n(\Theta) &= \sum_{n=1}^{N} \sum_{i=1}^{I} \log(\pi_i)\, \mathbb{I}\{\mathbf{z}_n^{\tau_{n1}} = i\} \\
&+ \sum_{n=1}^{N} \sum_{s=1}^{T_n-1} \sum_{i=1}^{I} \sum_{\substack{k=1 \\ k\neq i}}^{I} \left[\chi_{n\,ik}^{\tau_n(s)} \left(\log(\delta_{ik\,\mathbf{u}_n^{\tau_{ns}}}) + \rho'_{ik\,\mathbf{u}_n^{\tau_{ns}}} w_n \right) - \psi_{n\,i}^{\tau_n(s)}\, \delta_{ik\,\mathbf{u}_n^{\tau_{ns}}}\, \exp(\rho'_{ik\,\mathbf{u}_n^{\tau_{ns}}} w_n) \right] \\
&+ \sum_{n=1}^{N} \sum_{s=1}^{T_n} \sum_{i=1}^{I} \left[\log \binom{J-1}{\mathbf{y}_n^{\tau_{ns}}} + \mathbf{y}_n^{\tau_{ns}} \log(\mu_i) + (J-1-\mathbf{y}_n^{\tau_{ns}}) \log(1-\mu_i) \right] \mathbb{I}\{\mathbf{z}_n^{\tau_{ns}} = i\} \\
&+ \sum_{n=1}^{N} \sum_{s=1}^{T_n} \sum_{j=0}^{J-1} \left[\log \binom{L-1}{\mathbf{u}_n^{\tau_{ns}}} + \mathbf{u}_n^{\tau_{ns}} \log(\eta_j) + (L-1-\mathbf{u}_n^{\tau_{ns}}) \log(1-\eta_j) \right] \mathbb{I}\{\mathbf{y}_n^{\tau_{ns}} = j\}
\end{aligned}
$$

As discussed earlier, we need to take the expected value of the complete-data log likelihood function with respect to the posterior distribution of the underlying health state:

$$
\begin{aligned}
&\mathbb{E}[\log\, \mathcal{L}(\Theta) | \mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}] \\
&= \sum_{n=1}^{N} \sum_{i=1}^{I} \log(\pi_i)\, \mathbb{P}(\mathbf{z}_n^{\tau_{n1}} = i \,|\, \mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}) \\
&+ \sum_{n=1}^{N} \sum_{s=1}^{T_n-1} \sum_{i=1}^{I} \sum_{\substack{k=1 \\ k\neq i}}^{I} \sum_{l=0}^{L-1} \Big[\mathbb{E}\big[\chi_{n\,ik}^{\tau_n(s)} | \mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}\big] \big(\log(\delta_{ikl}) + \rho'_{ikl} w_n \big) \\
&\qquad - \mathbb{E}[\psi_{n\,i}^{\tau_n(s)} | \mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}]\, \delta_{ikl}\, \exp(\rho'_{ikl} w_n) \Big] \mathbb{I}\{\mathbf{u}_n^{\tau_{ns}} = l\} \\
&+ \sum_{n=1}^{N} \sum_{s=1}^{T_n} \sum_{i=1}^{I} \left[\log \binom{J-1}{\mathbf{y}_n^{\tau_{ns}}} + \mathbf{y}_n^{\tau_{ns}} \log(\mu_i) + (J-1-\mathbf{y}_n^{\tau_{ns}}) \log(1-\mu_i) \right] \\
&\qquad \times \mathbb{P}(\mathbf{z}_n^{\tau_{ns}} = i \,|\, \mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}) \\
&+ \sum_{n=1}^{N} \sum_{s=1}^{T_n} \sum_{j=0}^{J-1} \left[\log \binom{L-1}{\mathbf{u}_n^{\tau_{ns}}} + \mathbf{u}_n^{\tau_{ns}} \log(\eta_j) + (L-1-\mathbf{u}_n^{\tau_{ns}}) \log(1-\eta_j) \right] \mathbb{I}\{\mathbf{y}_n^{\tau_{ns}} = j\}
\end{aligned}
$$

In the E-step, we calculate the posterior probability of the underlying health state as well as the expected values of the end-state conditioned statistics. Afterward, in the M-step we can update the parameters by maximizing the expected log likelihood function. For this

purpose let $\gamma_n^s(i) = \mathbb{P}(\mathbf{z}_n^{\tau_n s} = i \,|\, \mathbf{u}_n^{\tau_n 1:T_n}, \mathbf{y}_n^{\tau_n 1:T_n}, \tilde{\Theta})$ (for $1 \le s \le T_n$) and $\nu_n^s(i,k) = \mathbb{P}(\mathbf{z}_n^{\tau_n s} = i, \mathbf{z}_n^{\tau_n s+1} = k \,|\, \mathbf{u}_n^{\tau_n 1:T_n}, \mathbf{y}_n^{\tau_n 1:T_n}, \tilde{\Theta})$ (for $1 \le s \le T_n - 1$) be the single-period and two-period posterior probabilities. To optimize over the initial state probability, set $\pi_I = 1 - \sum_{i'=1}^{I-1} \pi_{i'}$ and take the derivative to get:

$$\frac{\partial}{\partial \pi_i} \mathbb{E}\big[\log \mathcal{L}(\Theta) | \mathbf{u}_n^{\tau_n 1:T_n}, \mathbf{y}_n^{\tau_n 1:T_n}, \tilde{\Theta}\big] = \frac{\partial}{\partial \pi_i} \sum_{i'=1}^{I-1} \log(\pi_{i'}) \left(\sum_{n=1}^{N} \gamma_n^1(i') \right)$$
$$+ \frac{\partial}{\partial \pi_i} \log(1 - \sum_{i'=1}^{I-1} \pi_{i'}) \left(\sum_{n=1}^{N} \gamma_n^1(I) \right) = \frac{1}{\pi_i} \left(\sum_{n=1}^{N} \gamma_n^1(i) \right) - \frac{1}{\pi_I} \left(\sum_{n=1}^{N} \gamma_n^1(I) \right)$$

Hence, the optimal value for the parameter would be $\hat{\pi}_i = \frac{\sum_{n=1}^{N} \gamma_n^1(i)}{\sum_{i'=1}^{I} \sum_{n=1}^{N} \gamma_n^1(i')}$. For the emission parameter μ_i, we have:

$$\frac{\partial}{\partial \mu_i} \mathbb{E}\big[\log \mathcal{L}(\Theta) | \mathbf{u}_n^{\tau_n 1:T_n}, \mathbf{y}_n^{\tau_n 1:T_n}, \tilde{\Theta}\big] = \sum_{n=1}^{N} \sum_{s=1}^{T_n} \left[\frac{1}{\mu_i} \mathbf{y}_n^{\tau_n s} - \frac{1}{1-\mu_i} (J - 1 - \mathbf{y}_n^{\tau_n s}) \right] \gamma_n^s(i)$$
$$= \frac{1}{\mu_i} \left(\sum_{n=1}^{N} \sum_{s=1}^{T_n} \mathbf{y}_n^{\tau_n s} \gamma_n^s(i) \right) - \frac{1}{1-\mu_i} \left(\sum_{n=1}^{N} \sum_{s=1}^{T_n} (J - 1 - \mathbf{y}_n^{\tau_n s}) \gamma_n^s(i) \right)$$

which implies $\hat{\mu}_i = \frac{\sum_{n=1}^{N} \sum_{s=1}^{T_n} \mathbf{y}_n^{\tau_n s} \gamma_n^s(i)}{(J-1) \sum_{n=1}^{N} \sum_{s=1}^{T_n} \gamma_n^s(i)}$ (for $1 \le i \le I$). Similarly, for the intervention probabilities we have $\hat{\eta}_j = \frac{\sum_{n=1}^{N} \sum_{s=1}^{T_n} \mathbf{u}_n^{\tau_n s} \mathbb{I}\{\mathbf{y}_n^{\tau_n s} = j\}}{(L-1) \sum_{n=1}^{N} \sum_{s=1}^{T_n} \mathbb{I}\{\mathbf{y}_n^{\tau_n s} = j\}}$ (for $0 \le j \le J-1$). For the generator matrix parameter δ_{ikl}, the derivative would be:

$$\frac{\partial}{\partial \, \delta_{ikl}} \mathbb{E}\big[\log \mathcal{L}(\Theta) | \mathbf{u}_n^{\tau_n 1:T_n}, \mathbf{y}_n^{\tau_n 1:T_n}, \tilde{\Theta}\big] = \frac{1}{\delta_{ikl}} \sum_{n=1}^{N} \sum_{s=1}^{T_n - 1} \mathbb{E}\big[\chi_{n\,ik}^{\tau_n (s)} | \mathbf{u}_n^{\tau_n 1:T_n}, \mathbf{y}_n^{\tau_n 1:T_n}, \tilde{\Theta}\big] \mathbb{I}\{\mathbf{u}_n^{\tau_n s} = l\}$$
$$- \sum_{n=1}^{N} \sum_{s=1}^{T_n - 1} \mathbb{E}[\psi_{n\,i}^{\tau_n (s)} | \mathbf{u}_n^{\tau_n 1:T_n}, \mathbf{y}_n^{\tau_n 1:T_n}, \tilde{\Theta}] \, \exp(\rho_{ikl}^{'} \, w_n) \mathbb{I}\{\mathbf{u}_n^{\tau_n s} = l\}$$

which gives the following update rule:

$$\hat{\delta}_{ikl} = \frac{\sum_{n=1}^{N} \sum_{s=1}^{T_n - 1} \mathbb{E}\big[\chi_{n\,ik}^{\tau_n (s)} | \mathbf{u}_n^{\tau_n 1:T_n}, \mathbf{y}_n^{\tau_n 1:T_n}, \tilde{\Theta}\big] \mathbb{I}\{\mathbf{u}_n^{\tau_n s} = l\}}{\sum_{n=1}^{N} \sum_{s=1}^{T_n - 1} \mathbb{E}[\psi_{n\,i}^{\tau_n (s)} | \mathbf{u}_n^{\tau_n 1:T_n}, \mathbf{y}_n^{\tau_n 1:T_n}, \tilde{\Theta}] \, \exp(\hat{\rho}_{ikl}^{'} \, w_n) \mathbb{I}\{\mathbf{u}_n^{\tau_n s} = l\}} \tag{3.6}$$

Finally, we can use the Newton's method for optimizing the remaining parameter ρ_{ikl}:

$$\frac{\partial}{\partial \rho_{ikl}}\mathbb{E}[\log \mathcal{L}(\Theta)|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}] = \sum_{n=1}^{N}\sum_{s=1}^{T_n-1} \mathbb{I}\{\mathbf{u}_n^{\tau_{ns}} = l\}$$
$$\left(\mathbb{E}[\chi_{n\,ik}^{\tau_n\,(s)}|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta})] - \mathbb{E}[\psi_{n\,i}^{\tau_n\,(s)}|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}]\,\hat{\delta}_{ikl}\,\exp(\rho_{ikl}^{'} w_n)\right) w_n \quad (3.7)$$

$$\frac{\partial^2}{\partial \rho_{ikl}^2}\mathbb{E}[\log \mathcal{L}(\Theta)|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}] = -\sum_{n=1}^{N}\sum_{s=1}^{T_n-1} \mathbb{I}\{\mathbf{u}_n^{\tau_{ns}} = l\}$$
$$\left(\mathbb{E}[\psi_{n\,i}^{\tau_n\,(s)}|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}]\,\hat{\delta}_{ikl}\,\exp(\rho_{ikl}^{'} w_n)\right) w_n w_n^{'} \quad (3.8)$$

In particular, one can update $\hat{\delta}$ based on 3.6 and then update $\hat{\rho}$ according to 3.7 and 3.8.

Hence, the main challenge is in calculating the two aforementioned expectations as well as the two probabilities γ and ν. Following the latent-state conditioning idea that we discussed earlier, the expectations of the summary statistics can be represented in the following forms:

$$\mathbb{E}[\chi_{n\,ik}^{\tau_n\,(s)}|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta})] = \sum_{i'=1}^{I}\sum_{k'=1}^{I} \mathbb{E}[\chi_{n\,ik}^{\tau_n\,(s)}|\mathbf{z}_n^{\tau_{ns}} = i^{'}, \mathbf{z}_n^{\tau_n\,s+1} = k^{'}, \mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}]\,\nu_n^s(i', k')$$

$$\mathbb{E}[\psi_{n\,i}^{\tau_n\,(s)}|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}] = \sum_{i'=1}^{I}\sum_{k'=1}^{I} \mathbb{E}[\psi_{n\,i}^{\tau_n\,(s)}|\mathbf{z}_n^{\tau_{ns}} = i^{'}, \mathbf{z}_n^{\tau_n\,s+1} = k^{'}, \mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}]\,\nu_n^s(i', k')$$

Thus, the same method suggested by [LLL15] is applicable for calculating the expectations. Specifically, we evaluate the generator matrix using 3.3, 3.4 with $l = \mathbf{u}_n^{\tau_{ns}}$, and then represent the integrals based on 3.1 and 3.2, which allows for efficient calculation of them.

3.4.4 Calculating the Posterior Probabilities

As discussed earlier, the idea is that we construct a time-inhomogeneous DT-HMM based on the original CT-HMM and use the forward-backward algorithm for calculating the posterior probabilities of the latent variables. Specifically, denote by $\tilde{\mathbf{z}}$ the corresponding discrete-time latent variable in the DT-HMM that we construct, and let the transition probability matrix

of this model be $[\tilde{\mathcal{P}}_{n\,(s)}]_{ik} = \mathbb{P}(\tilde{\mathbf{z}}_n^{s+1} = k|\tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_{ns}}, \tilde{\Theta}) = \left[e^{(\tau_{n\,s+1}-\tau_{ns})\tilde{\mathcal{Q}}_{n\,(s)}}\right]_{ik}$.

Define $\alpha_n^s(i) = \mathbb{P}(\tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_{n\,1:s}}, \mathbf{y}_n^{\tau_{n\,1:s}}|\tilde{\Theta})$ and $\beta_n^s(i) = \mathbb{P}(\mathbf{u}_n^{\tau_{n\,s+1:T_n}}, \mathbf{y}_n^{\tau_{n\,s+1:T_n}}|\tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_{ns}}, \tilde{\Theta})$ as the forward and backward probabilities. In the following, we first present a set of dynamic programming equations for efficient calculation of α and β, and then show how γ and ν can be represented based on these quantities. Denote by $\tilde{\mu}$ and $\tilde{\eta}$ the emission and intervention parameters in the current EM iteration and suppose $\tilde{\mathcal{E}}(\cdot,\cdot)$ and $\tilde{\mathcal{G}}(\cdot,\cdot)$ indicate the corresponding model. Moreover, notice that the transition probabilities $\tilde{\mathcal{P}}_{n\,(s-1)} = e^{(\tau_{ns}-\tau_{n\,s-1})\tilde{Q}_{n\,(s-1)}}$ and $\tilde{\mathcal{P}}_{n\,(s)} = e^{(\tau_{n\,s+1}-\tau_{ns})\tilde{Q}_{n\,(s)}}$ are known since $\tilde{Q}_{n\,(s)}$ and $\tilde{Q}_{n\,(s-1)}$ are known for the current parameter values and the intervention choices. Hence, we can formulate the equations that govern the relationship between α variables and the β variables as:

$$
\begin{aligned}
\alpha_n^s(i) &= \sum_{k=1}^{I} \mathbb{P}(\tilde{\mathbf{z}}_n^s = i, \tilde{\mathbf{z}}_n^{s-1} = k, \mathbf{u}_n^{\tau_{n\,1:s}}, \mathbf{y}_n^{\tau_{n\,1:s}}|\tilde{\Theta}) \\
&= \sum_{k=1}^{I} \mathbb{P}(\mathbf{u}_n^{\tau_{ns}}|\mathbf{y}_n^{\tau_{ns}}, \tilde{\Theta})\mathbb{P}(\mathbf{y}_n^{\tau_{ns}}|\tilde{\mathbf{z}}_n^s = i, \tilde{\Theta})\,\mathbb{P}(\tilde{\mathbf{z}}_n^s = i|\tilde{\mathbf{z}}_n^{s-1} = k, \mathbf{u}_n^{\tau_{n\,s-1}}, \tilde{\Theta}) \\
&\qquad \mathbb{P}(\tilde{\mathbf{z}}_n^{s-1} = k, \mathbf{u}_n^{\tau_{n\,1:s-1}}, \mathbf{y}_n^{\tau_{n\,1:s-1}}|\tilde{\Theta}) \\
&= \sum_{k=1}^{I} \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}})\,\tilde{\mathcal{E}}(i, \mathbf{y}_n^{\tau_{ns}})\left[\tilde{\mathcal{P}}_{n\,(s-1)}\right]_{ki}\alpha_n^{s-1}(k) \qquad (3.9) \\
\beta_n^s(i) &= \sum_{k=1}^{I} \mathbb{P}(\tilde{\mathbf{z}}_n^{s+1} = k, \mathbf{u}_n^{\tau_{n\,s+1:T_n}}, \mathbf{y}_n^{\tau_{n\,s+1:T_n}}|\tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_{ns}}, \tilde{\Theta}) \\
&= \sum_{k=1}^{I} \mathbb{P}(\mathbf{u}_n^{\tau_{n\,s+2:T_n}}, \mathbf{y}_n^{\tau_{n\,s+2:T_n}}|\tilde{\mathbf{z}}_n^{s+1} = k, \mathbf{u}_n^{\tau_{n\,s+1}}, \tilde{\Theta})\,\mathbb{P}(\mathbf{u}_n^{\tau_{n\,s+1}}|\mathbf{y}_n^{\tau_{n\,s+1}}, \tilde{\Theta}) \\
&\qquad \mathbb{P}(\mathbf{y}_n^{\tau_{n\,s+1}}|\tilde{\mathbf{z}}_n^{s+1} = k, \tilde{\Theta})\,\mathbb{P}(\tilde{\mathbf{z}}_n^{s+1} = k|\tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_{ns}}, \tilde{\Theta}) \\
&= \sum_{k=1}^{I} \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_{n\,s+1}}, \mathbf{u}_n^{\tau_{n\,s+1}})\,\tilde{\mathcal{E}}(k, \mathbf{y}_n^{\tau_{n\,s+1}})\left[\tilde{\mathcal{P}}_{n\,(s)}\right]_{ik}\beta_n^{s+1}(k) \qquad (3.10)
\end{aligned}
$$

The corner cases can also be determined by:

$$
\alpha_n^1(i) = \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_{n1}}, \mathbf{u}_n^{\tau_{n1}})\,\tilde{\mathcal{E}}(i, \mathbf{y}_n^{\tau_{n1}})\tilde{\pi}_i \qquad (3.11)
$$

$$\beta_n^{T_n-1}(i) = \sum_{k=1}^{I} \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_n T_n}, \mathbf{u}_n^{\tau_n T_n})\, \tilde{\mathcal{E}}(k, \mathbf{y}_n^{\tau_n T_n})\, \big[\tilde{\mathcal{P}}_{n\,(T_n-1)}\big]_{ik} \tag{3.12}$$

Now, one can easily show that $\gamma_n^s(i) = \frac{\alpha_n^s(i)\beta_n^s(i)}{\sum_{i'=1}^{I}\alpha_n^s(i')\beta_n^s(i')}$. Furthermore, notice that:

$$\mathbb{P}(\tilde{\mathbf{z}}_n^s = i, \tilde{\mathbf{z}}_n^{s+1} = k, \mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}|\tilde{\Theta}) = \mathbb{P}(\mathbf{u}_n^{\tau_n\,s+2:T_n}, \mathbf{y}_n^{\tau_n\,s+2:T_n}|\tilde{\mathbf{z}}_n^{s+1} = k, \mathbf{u}_n^{\tau_n\,s+1}, \tilde{\Theta}) \times$$
$$\mathbb{P}(\mathbf{u}_n^{\tau_n\,s+1}|\mathbf{y}_n^{\tau_n\,s+1}, \tilde{\Theta})\, \mathbb{P}(\mathbf{y}_n^{\tau_n\,s+1}|\tilde{\mathbf{z}}_n^{s+1} = k, \tilde{\Theta})\, \mathbb{P}(\tilde{\mathbf{z}}_n^{s+1} = k|\tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_n\,s}, \tilde{\Theta})\, \mathbb{P}(\tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_n\,1:s}, \mathbf{y}_n^{\tau_n\,1:s}, \tilde{\Theta})$$
$$= \beta_n^{s+1}(k)\tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_n\,s+1}, \mathbf{u}_n^{\tau_n\,s+1})\, \tilde{\mathcal{E}}(k, \mathbf{y}_n^{\tau_n\,s+1})\big[\tilde{\mathcal{P}}_{n\,(s)}\big]_{ik}\, \alpha_n^s(i)$$
$$\mathbb{P}(\tilde{\mathbf{z}}_n^{T_n-1} = i, \tilde{\mathbf{z}}_n^{T_n} = k, \mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}|\tilde{\Theta}) = \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_n T_n}, \mathbf{u}_n^{\tau_n T_n})\, \tilde{\mathcal{E}}(k, \mathbf{y}_n^{\tau_n T_n})\big[\tilde{\mathcal{P}}_{n\,(T_n-1)}\big]_{ik}\, \alpha_n^{T_n-1}(i)$$

Hence, the remaining probability can be represented based on α and β as:

$$\nu_n^s(i,k) = \frac{\tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_n\,s+1}, \mathbf{u}_n^{\tau_n\,s+1})\, \tilde{\mathcal{E}}(k, \mathbf{y}_n^{\tau_n\,s+1})\big[\tilde{\mathcal{P}}_{n\,(s)}\big]_{ik}\, \alpha_n^s(i)\, \beta_n^{s+1}(k)}{\sum_{i'=1}^{I}\sum_{k'=1}^{I}\tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_n\,s+1}, \mathbf{u}_n^{\tau_n\,s+1})\, \tilde{\mathcal{E}}(k', \mathbf{y}_n^{\tau_n\,s+1})\big[\tilde{\mathcal{P}}_{n\,(s)}\big]_{i'\,k'}\, \alpha_n^s(i')\, \beta_n^{s+1}(k')} \tag{3.13}$$

And the corner case would also be similar.

3.4.5 Numerical Experiments

To evaluate the quality of the discussed estimation algorithm, we perform numerical experiments on synthetically-generated data. The algorithm has been implemented in *R* and visualizations have been performed in *Python*. Please refer to *GitHub*[1] for all the codes and results.

At first, we briefly explain the procedure that has been used for generating the samples. For each patient, the number of visits is a random draw from a Gaussian distribution as $T_n \sim round(\mathcal{N}(\mu_T, \sigma_T^2))$. The inter-visit times are also determined based on a Gaussian distribution $(\tau_{n\,s+1} - \tau_{n\,s}) \sim \mathcal{N}(\mu_\tau, \sigma_\tau^2)$ for $1 \leq s \leq T_n$ (we assume $\tau_{n1} = 0$). In the simulations, we set $I = 3$, $J = 10$, and $L = 3$. We draw the initial state from a categorical distribution $\mathbf{z}_n^{\tau_{n1}} \sim Cat(\pi^*)$ that reflects our belief regarding the distribution of the underlying

[1] Refer to the CT-HMM directory in *https://github.com/saeedghodsi93/Disease_Progression_Modeling_HMM*

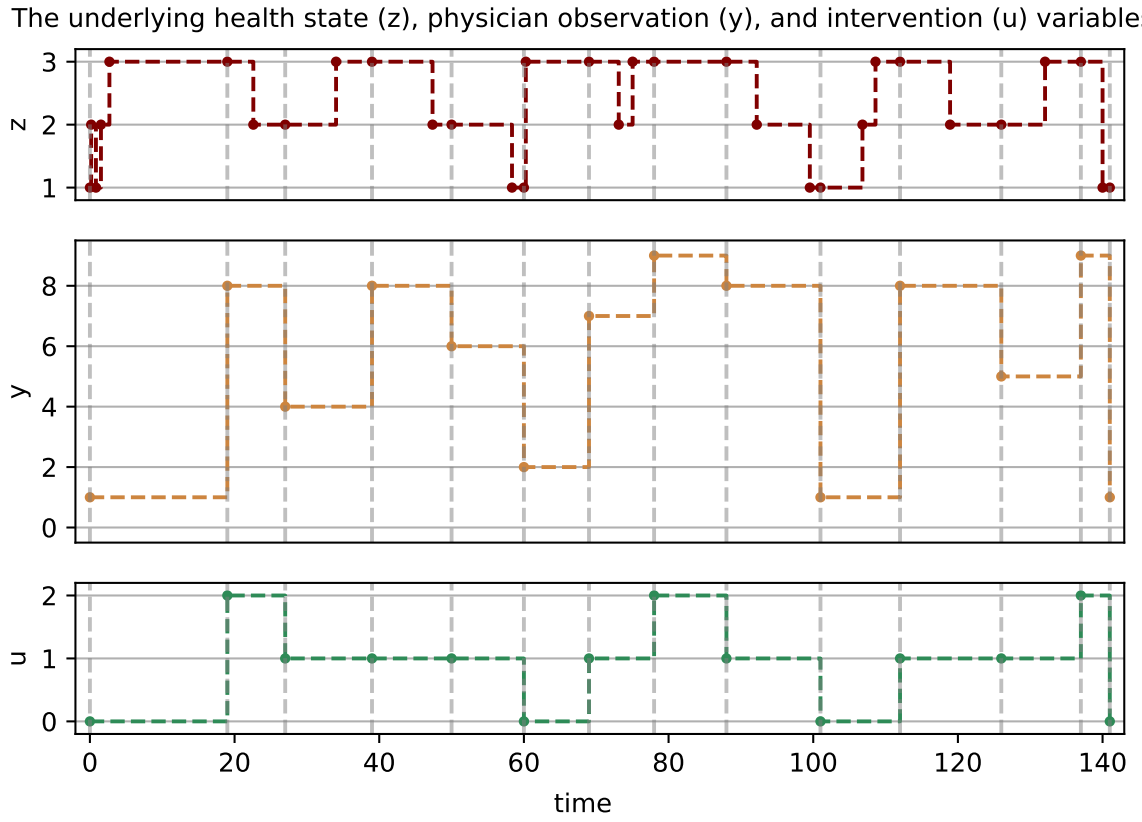

Figure 3.2: Synthetically-generated trajectories.

variable across the population. Similarly, the emission and intervention distributions are categorical as $\mathbf{y}_n^{\tau_{n1}} \sim Cat(\mathcal{E}^*(\mathbf{z}_n^{\tau_{n1}}, \cdot))$ and $\mathbf{u}_n^{\tau_{n1}} \sim Cat(\mathcal{G}^*(\mathbf{y}_n^{\tau_{n1}}, \cdot))$. Age is also used as the only covariate, and it is generated according to the mixture of two Gaussian distributions as $a_n \sim p_a \mathcal{N}(\mu_{a,y}, \sigma_{a,y}^2) + (1 - p_a)\mathcal{N}(\mu_{a,o}, \sigma_{a,o}^2)$.

The remaining tricky part is generating the underlying health state variable. Suppose ζ^* and ι^* are the transition matrix and the mean sojourn time parameters of the embedded DTMC. For any time interval s, we calculate these parameters (i.e. $\zeta^*_{\cdot,\cdot,\mathbf{u}_n^{\tau_{ns}}}$ and $\iota^*_{\cdot,\mathbf{u}_n^{\tau_{ns}}}$) using $\mathcal{Q}^*_{\cdot,\cdot,\mathbf{u}_n^{\tau_{ns}}}$. In a given state, we draw the sojourn time from $(\tilde{\tau}_{n\,\tilde{s}+1} - \tilde{\tau}_{n\tilde{s}}) \sim Exp(\iota^*_{\mathbf{z}_n^{\tilde{\tau}_{n\tilde{s}}},\mathbf{u}_n^{\tau_{ns}}})$ and the destination state from $\mathbf{z}_n^{\tilde{\tau}_{n\,\tilde{s}+1}} \sim Cat(\zeta^*_{\mathbf{z}_n^{\tilde{\tau}_{n\tilde{s}}},\cdot,\mathbf{u}_n^{\tau_{ns}}})$. Afterward, we record the transition information and repeat the same procedure for the next move. The iterations stop whenever $\tilde{\tau}_{n\,\tilde{s}+1} \geq \tau_{ns}$ for some $\tilde{s}$. At the observation time, we generate the emission and intervention

variables as discussed earlier.

The next step would then be to run the EM algorithm on the generated dataset and estimate the parameters. To evaluate the performance of the algorithm, one can measure the Root Mean Squared Error (RMSE) of the estimated parameter with respect to the true parameter. Figure 3.3 plots the convergence of the parameters to the true parameters as measured by RMSE. We observe that the algorithm is able to reconstruct the true parameters givne enough number of samples and enough number of EM iterations.

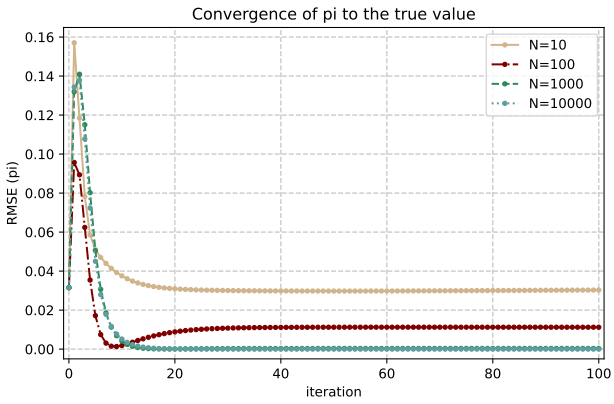

(a) Convergence of $\hat{\pi}$

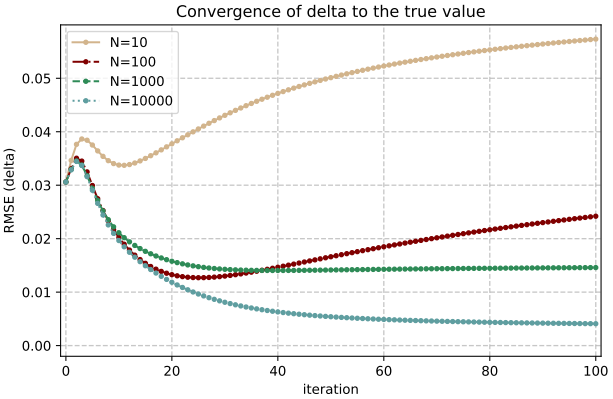

(b) Convergence of $\hat{\delta}$

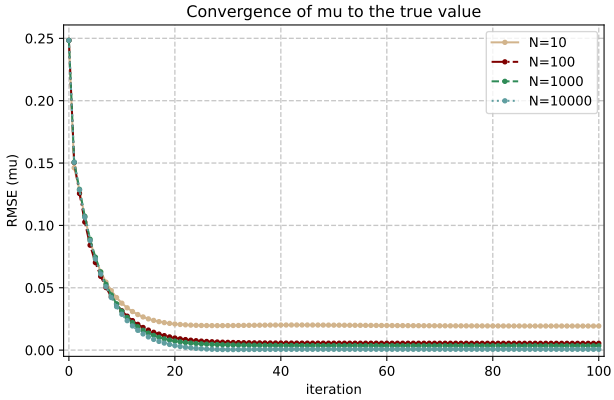

(c) Convergence of $\hat{\mu}$

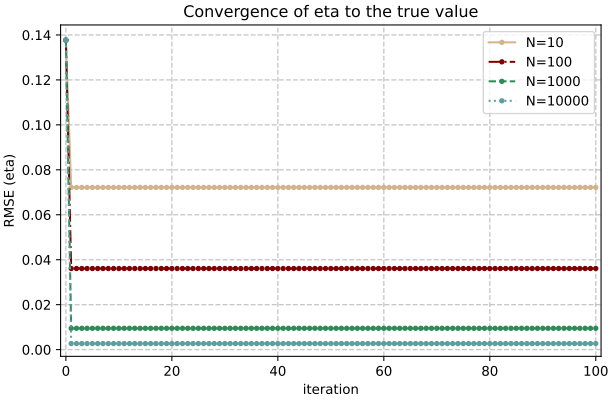

(d) Convergence of $\hat{\eta}$

Figure 3.3: The RMSE plots for the estimated parameters.

3.5 Population-Level Planning

Suppose we have a population of patients under support (e.g. through a hybrid FFS-capitation healthcare payment system). The healthcare provider aims to keep the patients as healthy as possible, under a set of resource and budget constraints. For this purpose, we employ the learned disease progression model and formulate a planning problem that is supposed to inform decisions regarding the choice of healthcare intervention strategies. In our framework, the flow of patients is described using a population-level model. More specifically, we assume the service needs of patients depends on the physician observation. Therefore, we divide the patients into multiple groups based on their physician observation values, and analyze the dynamics of the group populations based on the chosen intervention distributions. The patient flow model that we propose is a linear model with deterministic stationary transition rates which are estimated based on the learned disease progression model. We believe the transitions of patients across different health states can be described appropriately using a linear model and a potential source of non-linearity in our model should be incorporated in resource consumption or budget constraints rather than the transition rates. Furthermore, such a linear model makes it more straightforward to interpret the results. The stationarity assumption also seems justifiable, since most probably the characteristics of the population will not change drastically as long as the duration of the planning horizon is not too long. Moreover, we solve the optimization problem again at the beginning of each time period, as we'll discuss later.

We model the decision problem as a multi-period planning problem and denote by $s = 0, 1, \cdots, T$ the time index. Suppose there are a set of distinct packages of healthcare interventions (including the no-intervention option) and each patient is allowed to receive at most one of the intervention categories at any given time period. At time point $s = 0$, the healthcare provider has access to the distribution of physician observations across the population under support, and decides about the number of patients to which each intervention

package is assigned for each period $1, \cdots, T$. Ideally, the provider would like to observe the underlying health conditions and choose the intervention on that basis. However, the most accurate available information is the initial physician observation data, and planning is performed using this proxy. Hence, we assume that the cost and resource consumption that are associated with providing interventions are determined based on the physician observation variable rather than the underlying health condition. We use a Model Predictive Control (MPC) optimization approach in which at the beginning of each period the planning is performed for a horizon of length T given the existing information, and it is solved again at the beginning of the next period when new information becomes available.

We model the provider's objective as a linear function of the number of patients in different physician observation categories. Specifically, we assume that patients in a given category j may be referred by their physician to receive a non-hospitalization service during any given period of time at the rate of $\varkappa_j^{NHSP}$. Similarly, we denote by $\varkappa_j^{HSP}$ the rate at which a patient in category j may receive a hospitalization service during any arbitrary period of time. Notice that $\varkappa_j^{NHSP}$ and $\varkappa_j^{HSP}$ may be smaller or greater than one depending on many factors including the duration of each time period. Let's assume a non-hospitalization service and a hospitalization service incur a fixed cost of S_j^{NHSP} and S_j^{HSP} to the system, respectively. Again, we're assuming the costs of these services for a given patient depend on physician observation rather than the underlying health condition. As explained beforehand, we also assume that payment to the provider is performed based on a hybrid FFS-capitation system. More specifically, any non-hospitalization service provided to the patients is compensated according to an FFS payment system with a coefficient Υ. Therefore, for any non-hospitalization service that a patient in category j receives, the provider incurs a cost equal to S_j^{NHSP} and gets reimbursed by an amount equal to $\Upsilon \times S_j^{NHSP}$. Following the literature, we assume $\Upsilon \leq 1$. Define the average net cost per patient per period as $\beta_j = (1-\Upsilon)\varkappa_j^{NHSP} S_j^{NHSP} + \varkappa_j^{HSP} S_j^{HSP}$ $(0 \leq j \leq J-1)$. It seems safe to assume that $\varkappa_j^{NHSP}$, $\varkappa_j^{HSP}$, S_j^{NHSP}, and S_j^{HSP} are non-decreasing in j or equivalently, β_j is non-decreasing in j. Moreover, an additional lump sum

$\bar{Q}$ is paid for both the hospitalization and non-hospitalization services needed by each patient during each period of time.

Let P_j^s $(0 \leq j \leq J-1,\, 0 \leq s \leq T)$ indicate the population in category j at time point s. Without loss of generality, we assume the population is normalized in a way that $\sum_{j=0}^{J-1} P_j^s = 1$. The provider's net payoff per patient during the entire planning horizon, denoted by ω, can be represented in the following way in terms of the distribution of population across health states over time:

$$\omega = \sum_{s=1}^{T}(\bar{Q} - \sum_{j=0}^{J-1} \beta_j\, P_j^s) = T\,\bar{Q} - \sum_{s=1}^{T}\sum_{j=0}^{J-1} \beta_j\, P_j^s$$

Notice that the initial population distribution is assumed to be given. Clearly, the provider's decisions do not affect the first term. Hence, maximizing the net payoff is equivalent to minimizing the second term above. We represent the objective function of the provider's problem as $\mathcal{F} = \sum_{s=1}^{T}\sum_{j=0}^{J-1} \beta_j P_j^s$. In other words, the provider is penalized in a linear way as the distribution of population's health state moves towards getting worse.

Denote by X_{jl}^s $(0 \leq j \leq J-1,\, 0 \leq l \leq L-1,\, 1 \leq s \leq T)$, the number of patients with physician observation j who receive intervention l in time period $[s-1, s)$ (so $\sum_{l=0}^{L-1} X_{jl}^s = P_j^{s-1}$). The total population in category j at time point s is composed of the set of patients who transitioned to category j from all the other categories during time period $[s-1, s)$ as well as patients who have remained in this category. Therefore, the flow of patients between health states across the time can be represented in the following way:

$$P_{j'}^s = \sum_{j=0}^{J-1}\sum_{l=0}^{L-1} U_{jl}^{j'} X_{jl}^s; \quad (\forall\, 0 \leq j' \leq J-1,\, 1 \leq s \leq T)$$

where $U_{jl}^{j'}$ represents the transition rate from category j to category j' under intervention option l. More specifically, $U_{jl}^{j'}$ is defied as the ratio of number of patients in category j who received intervention option l and transitioned to category j' after one period of time divided by the total number of patients in category j who received intervention option l in that period

of time. We assume this ratio is deterministic and constant over time. Clearly, the transition rates matrix should satisfy $\sum_{j'=0}^{J-1} U_{jl}^{j'} = 1$ for any any $0 \leq j \leq J-1$, $0 \leq l \leq L-1$. For the moment, assume the rates are known. We'll discuss the estimation approach later in this section.

The objective function of the problem can be simplified as:

$$\mathcal{F}(X) = \sum_{s=1}^{T} \sum_{j'=0}^{J-1} \beta_{j'} P_{j'}^{s} = \sum_{s=1}^{T} \sum_{j'=0}^{J-1} \beta_{j'} (\sum_{j=0}^{J-1} \sum_{l=0}^{L-1} U_{jl}^{j'} X_{jl}^{s}) = \sum_{s=1}^{T} \sum_{j=0}^{J-1} \sum_{l=0}^{L-1} X_{jl}^{s} (\sum_{j'=0}^{J-1} \beta_{j'} U_{jl}^{j'}) \quad (3.14)$$

We can interpret the term $\sum_{j'=0}^{J-1} \beta_{j'} U_{jl}^{j'}$ as the total expected contribution of a patient in group j to the objective function during an arbitrary period of time, after receiving intervention option l. Furthermore, the patient flow constraints can be simplified in the following way:

$$\sum_{l=0}^{L-1} X_{j'l}^{1} = P_{j'}^{0}; \quad (\forall\, 0 \leq j' \leq J-1) \quad (3.15)$$

$$\sum_{l=0}^{L-1} X_{j'l}^{s} = P_{j'}^{s-1} = \sum_{j=0}^{J-1} \sum_{l=0}^{L-1} U_{jl}^{j'} X_{jl}^{s-1}; \quad (\forall\ 0 \leq j' \leq J-1,\ 2 \leq s \leq T) \quad (3.16)$$

We assume there are M resources in the system and $\bar{R}_m^s$ $(1 \leq m \leq M,\, 1 \leq s \leq T)$ indicates the total amount of available resource m in time period $[s-1, s)$. Moreover, let's denote by R_{jl}^{m} $(0 \leq j \leq J-1,\, 0 \leq l \leq L-1,\, 1 \leq m \leq M)$ the amount of resource m required for providing intervention l to a patient in category j for one period of time. Similarly, we assume that there is a certain amount of available budget $\bar{B}$ for the entire planning horizon and providing intervention l to a patient in category j incurs a fixed cost C_{jl}. Furthermore, some intervention options may be naturally unsuitable for certain classes of health states for any reason. Consider all the pairs of indices (j, l) $(0 \leq j \leq J-1,\, 0 \leq l \leq L-1)$ for which we know that intervention l must be avoided for patients in category j. Define a hypothetical resource consumption matrix indexed by $m = 0$ and set all the elements corresponding to these pairs equal to 1. Also, fill all the other elements with zeros and set the amount of

available resource $m = 0$ equal to zero in all the periods of time, i.e. $\bar{R}_0^s = 0$ for all $1 \leq s \leq T$. By adding constraints that guarantee no resource $m = 0$ is being consumed at any given period of time, we're practically guaranteeing that the decision variables X_{jl}^s will be set to zero for all of those pairs (j, l). The following resource consumption and budget constraints then must be satisfied:

$$\sum_{j=0}^{J-1} \sum_{l=0}^{L-1} X_{jl}^s R_{jl}^m \leq \bar{R}_m^s; \quad (\forall\ 0 \leq m \leq M,\ 1 \leq s \leq T) \tag{3.17}$$

$$\sum_{s=1}^{T} \sum_{j=0}^{J-1} \sum_{l=0}^{L-1} X_{jl}^s C_{jl} \leq \bar{B} \tag{3.18}$$

Notice that in our framework we do not explicitly model how many times an intervention needs to be applied to patients during a given period of time. For example, a patient suffering from major depression may need multiple psychotherapy sessions during one period of time, while one session per period might be sufficient for a patient with early-stage depression. The impact of this heterogeneity in intensity and frequency of required interventions is reflected through resource consumption and intervention cost parameters. Specifically, for some intervention options the resource consumption parameters R_{jl}^m are increasing in j. On the other hand, the consumption of certain resources by some interventions is almost independent from the disease severity. For example, a hospitalized patient requires a single bed while the need for nursing staff may depend on the severity of the disease.

Finally, all the decision variables must have non-negative values:

$$X_{jl}^s \geq 0; \quad (\forall 0 \leq j \leq J-1,\ 0 \leq l \leq L-1, 1 \leq s \leq T) \tag{3.19}$$

The optimization problem can then be formulated as:

$$\min_{X} \quad \mathcal{F}(X) = \sum_{s=1}^{T} \sum_{j=0}^{J-1} \sum_{l=0}^{L-1} X_{jl}^{s} (\sum_{j'=0}^{J-1} \beta_{j'} U_{jl}^{j'})$$

$$s.t. \quad 3.15, 3.16, 3.17, 3.18, 3.19$$

From the perspective of policy implementation, we prefer 0/1 solutions, since they make application of the policy less challenging. In particular, implementing a fractional solution requires assigning different interventions to patients with the same characteristics, which may raise fairness issues. One way to resolve this issue would be to add integrality constraints to the problem, which can potentially lead to making sub-optimal decisions. Although the solution of the planning problem without integrality constraints is not guaranteed to be integer in general, we expect the solution to be integer for the above problem, as it is an LP.

The only remaining part is to estimate the matrix U. To estimate the transition rates, we use the learned disease progression model. Suppose we have estimated the parameter set $\hat{\Theta} = (\hat{\pi}, \hat{\rho}, \hat{\delta}, \hat{\mu}, \hat{\eta})$ based on historical data. Moreover, assume the duration of each time period is $\Delta\tau$. We model the flow of the population between different physician observation categories based on the transitions of a representative patient. Let $\tilde{\mathbf{z}}^0$ and $\tilde{\mathbf{y}}^0$ denote the underlying health state and the physician observation variable, respectively, at the beginning of planning. Similarly, define $\tilde{\mathbf{z}}^1$ and $\tilde{\mathbf{y}}^1$ as the corresponding variables at the end of the first period. Furthermore, suppose $\tilde{\mathbf{u}}^1$ is the intervention variable during the first period. As explained earlier, we assume the initial distribution of $\tilde{\mathbf{y}}$ is known, and plan for the entire horizon $s = 1, \cdots, T$. To do so, we approximate the transition rate between different categories of $\tilde{\mathbf{y}}$ for all the time periods with the transition rate of the first period. By minimizing the expected objective function under the resource and budget constraints, we find the optimal distribution of intervention assignments for each of the T time periods. However, at the execution time we only execute the plan for the first period and plan again at the beginning of the next period when the updated information on $\tilde{\mathbf{y}}$ is available. As a consequence, the

transition rates matrix will also be estimated again in the next execution step.

If the initial physician observation is $\tilde{\mathbf{y}}^0 = j$, the posterior probability of the initial underlying health state will be $\mathbb{P}(\tilde{\mathbf{z}}^0 = i \,|\, \tilde{\mathbf{y}}^0 = j) = \frac{\hat{\mathcal{E}}(i,j)\,\hat{\pi}_i}{\sum_{i'=1}^{I} \hat{\mathcal{E}}(i',j)\,\hat{\pi}_{i'}}$. Suppose the representative patient is in state i at time point $s = 0$. The probability $\mathbb{P}(\tilde{\mathbf{z}}^1 = k \,|\, \tilde{\mathbf{z}}^0 = i, \tilde{\mathbf{u}}^1 = l)$ of transition to state k in time point 1 under intervention $\tilde{\mathbf{u}}^1 = l$ can then be estimated as $\left[e^{\Delta\tau\,\hat{\mathcal{Q}}_l}\right]_{ik}$ where $\hat{\mathcal{Q}}_l$ is the estimated generator matrix for intervention l (notice that here we're assuming the population is homogeneous). Hence, the posterior distribution of the underlying health state at time point $s = 1$ can then be expressed as:

$$\begin{aligned}\mathbb{P}(\tilde{\mathbf{z}}^1 = k \,|\, \tilde{\mathbf{y}}^0 = j, \tilde{\mathbf{u}}^1 = l) &= \sum_{i=1}^{I} \mathbb{P}(\tilde{\mathbf{z}}^0 = i, \tilde{\mathbf{z}}^1 = k \,|\, \tilde{\mathbf{y}}^0 = j, \tilde{\mathbf{u}}^1 = l) \\ = \sum_{i=1}^{I} \mathbb{P}(\tilde{\mathbf{z}}^1 = k \,|\, \tilde{\mathbf{z}}^0 = i, \tilde{\mathbf{u}}^1 = l)\, \mathbb{P}(\tilde{\mathbf{z}}^0 = i \,|\, \tilde{\mathbf{y}}^0 = j) &= \frac{\sum_{i=1}^{I} \hat{\mathcal{E}}(i,j)\,\hat{\pi}_i \left[e^{\Delta\tau\,\hat{\mathcal{Q}}_l}\right]_{ik}}{\sum_{i'=1}^{I} \hat{\mathcal{E}}(i',j)\,\hat{\pi}_{i'}}\end{aligned}$$

By aggregating over all the possible values for j and l, we can calculate the distribution of $\tilde{\mathbf{z}}^1$:

$$\begin{aligned}\mathbb{P}(\tilde{\mathbf{z}}^1 = k) &= \sum_{j=0}^{J-1} \sum_{l=0}^{L-1} \mathbb{P}(\tilde{\mathbf{z}}^1 = k \,|\, \tilde{\mathbf{y}}^0 = j, \tilde{\mathbf{u}}^1 = l)\, \mathbb{P}(\tilde{\mathbf{y}}^0 = j, \tilde{\mathbf{u}}^1 = l) \\ &= \sum_{j=0}^{J-1} \sum_{l=0}^{L-1} \frac{\sum_{i=1}^{I} \hat{\mathcal{E}}(i,j)\,\hat{\pi}_i \left[e^{\Delta\tau\,\hat{\mathcal{Q}}_l}\right]_{ik}}{\sum_{i'=1}^{I} \hat{\mathcal{E}}(i',j)\,\hat{\pi}_{i'}}\, X_{jl}^1 = \sum_{j=0}^{J-1} \sum_{l=0}^{L-1} \varrho_{jl}^k\, X_{jl}^1 \qquad (3.20)\end{aligned}$$

where $\varrho_{jl}^k = \frac{\sum_{i=1}^{I} \hat{\mathcal{E}}(i,j)\,\hat{\pi}_i \left[e^{\Delta\tau\,\hat{\mathcal{Q}}_l}\right]_{ik}}{\sum_{i'=1}^{I} \hat{\mathcal{E}}(i',j)\,\hat{\pi}_{i'}}$ is known for given values of j, l, and k. Notice that $\sum_{k=1}^{I} \varrho_{jl}^k = 1$, which implies that $\sum_{k=1}^{I} \mathbb{P}(\tilde{\mathbf{z}}^1 = k) = 1$, as expected. The distribution of physician observation at time point $s = 1$ can then be calculated as:

$$\begin{aligned}P_{j'}^1 = \mathbb{P}(\tilde{\mathbf{y}}^1 = j') &= \sum_{k=1}^{I} \mathbb{P}(\tilde{\mathbf{y}}^1 = j' \,|\, \tilde{\mathbf{z}}^1 = k)\, \mathbb{P}(\tilde{\mathbf{z}}^1 = k) = \sum_{k=1}^{I} \hat{\mathcal{E}}(k,j') \left(\sum_{j=0}^{J-1} \sum_{l=0}^{L-1} \varrho_{jl}^k\, X_{jl}^1 \right) \\ &= \sum_{j=0}^{J-1} \sum_{l=0}^{L-1} \left(\sum_{k=1}^{I} \hat{\mathcal{E}}(k,j')\, \varrho_{jl}^k \right) X_{jl}^1 = \sum_{j=0}^{J-1} \sum_{l=0}^{L-1} U_{jl}^{j'}\, X_{jl}^1 \qquad (3.21)\end{aligned}$$

where the transition rate $U_{jl}^{j^{'}} = \sum_{k=1}^{I} \hat{\mathcal{E}}(k, j^{'})\, \varrho_{jl}^{k} = \frac{\sum_{i=1}^{I} \hat{\mathcal{E}}(i,j)\, \hat{\pi}_i \left(\sum_{k=1}^{I} \hat{\mathcal{E}}(k,j^{'}) \left[e^{\Delta\tau\, \hat{\mathcal{Q}}_l} \right]_{ik} \right)}{\sum_{i^{'}=1}^{I} \hat{\mathcal{E}}(i^{'},j)\, \hat{\pi}_{i^{'}}}$ is known for given values of j, l, and $j^{'}$. As Equations 3.20 and 3.21 indicate, both the distributions of the underlying health state and the physician observation can be represented linearly in terms of the decision variable X_{jl}^{1}.

3.6 Model Extension: Occasionally-Observed Underlying Health Conditions

In this section, we study some other healthcare domains that have different characteristics than those behavioral conditions that we've been considering so far. Specifically, we develop an alternative disease progression model in which the underlying health condition is occasionally revealed through accurate examination. For example, echocardiogram, which is a type of ultrasound imaging test, helps physicians to analyze the structure and function of heart in patients with potential heart disease. This information is much more detailed than the results of typical examinations, lab tests, and blood pressure measurement. Therefore, it is reasonable to assume that it gives a measure of the actual underlying health condition of the patient. Similarly, biopsy as well as a variety of imaging tests such as Computerized Tomography (CT) scan, Magnetic Resonance Imaging (MRI), Positron Emission Tomography (PET) scan, and X-ray are used as tools for cancer diagnosis. However, accurate examination is often costly and imposes some risks to the patient. Hence, such detailed examinations occur only when physician observations that are collected during regular tests and visits suggest that the patient's health condition is getting worse. In this section, we modify the baseline disease progression model by assuming that the underlying health condition variable is occasionally observed by accurate examination.

Although our model has few restrictive assumptions and can potentially fit to a variety of contexts with minor adjustments, we consider maintenance therapy for cancer as our main medical domain. Maintenance therapy targets patients in the remission phase, and tries to prevent the cancer's return or delay the growth of advanced cancer after the initial treatment.

Hormonal therapy, immunotherapy, and chemotherapy are the most important categories of interventions that are commonly used during maintenance therapy. Besides the choice of intervention type, in this context it is crucial to accurately tune the dosage of the medication or aggressiveness of the tissue removal. In particular, prescribing a high dosage of medication or removing too much of the body tissues may result in side effects on other body organs.

In addition to the above modification, we assume the chosen intervention at any given period has a direct effect on the physician observation in the next period. In other words, the intervention may also affect the physician observation through a separate direct path in addition to affecting the underlying course of disease progression. The rationale behind this idea is that in some certain healthcare domains interventions have negligible impact on slowing or reversing the progression trend of the disease. However, they may serve as a symptoms-control mechanism, which means that they have a direct effect on what the physician observes. For example, symptom-control medications that are sometimes prescribed for patients with terminal illness mainly affect the physician observations rather than the actual course of disease progression. Therefore, in our extended model we assume that there is a link from the intervention variable of each time period to the corresponding physician observation variable of the next time period. Notice that the strength of the relationship between intervention and physician observation may decay over time due to issues related to medication adherence in cases in which there is a long between-visit time. Although the amount of time between consecutive visits can potentially vary in our model, for the sake of mathematical simplicity here we assume that the effect of intervention on physician observation is constant across all the periods.

3.6.1 Model Setup

For patient n, denote by $\mathcal{O}_n \subseteq \{1, \cdots, T_n\}$ the set of indices that correspond to regular visit times. Moreover, let $\mathcal{O}_n'$ be the complement of $\mathcal{O}_n$, i.e. the set of indices that correspond to accurate examination. By definition, the variable z is unobserved during regular visits,

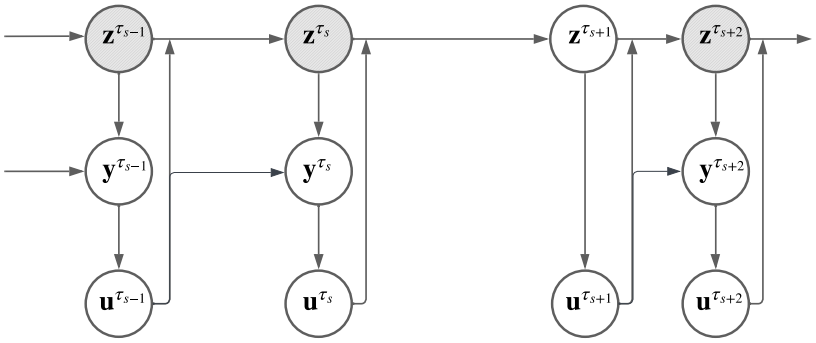

Figure 3.4: A modified version of the CT-HMM disease progression model in which the z variables are sometimes directly observed and used for choosing the intervention (shadowed nodes indicate unobserved variables). Moreover, the chosen intervention at any given time period has a direct effect on the physician observation in the next period.

but observed in accurate examinations. We consider two decision making mechanisms that correspond to these two types of visits. Specifically, binomial distributions with parameters η_j $(0 \leq j \leq J-1)$ map the physician observations to the intervention choice in regular visits, similar to the baseline model. On the other hand, separate binomial distributions with parameters $\eta_i^{'}$ $(1 \leq i \leq I)$ map the distribution of underlying health condition to the interventions in accurate-examination periods. Also, suppose the direct intervention distribution is denoted by $\mathcal{G}^{'}(i,l) = \mathbb{P}(\mathbf{u}_n^{\tau_{ns}} = l \,|\, \mathbf{z}_n^{\tau_{ns}} = i)$. In Figure 3.4, we present the graphical model associated with the described scenario. As the figure illustrates, variable $\mathbf{y}_n^{\tau_{ns}}$ is affected by both $\mathbf{z}_n^{\tau_{ns}}$ and $\mathbf{u}_n^{\tau_n s-1}$. Hence, we model $\mathbf{y}_n^{\tau_{ns}}$ as a binomial distribution with parameter μ_{il} for $\mathbf{z}_n^{\tau_{ns}} = i$ $(1 \leq i \leq I)$ and $\mathbf{u}_n^{\tau_n s-1} = l$ $(0 \leq l \leq L-1)$. We also denote the emission distribution by $\mathcal{E}(i,l,j) = \mathbb{P}(\mathbf{y}_n^{\tau_{ns}} = j \,|\, \mathbf{z}_n^{\tau_{ns}} = i, \mathbf{u}_n^{\tau_n s-1} = l)$. Needless to explain, there is no variable $\mathbf{y}_n^{\tau_{ns}}$ in periods $s \in \mathcal{O}_n^{'}$.

3.6.2 Parameter Learning via the EM Algorithm

The set of parameters that we're interested to estimate is $\Theta = (\pi, \rho, \delta, \mu, \eta, \eta^{'})$. The complete-data log-likelihood can be expressed in the following way, similar to Equation 3.5:

$$\log\, \mathcal{L}_n(\Theta) = \log \mathbb{P}(\mathbf{z}_n^{\tau_{n1}}) + \sum_{s=1}^{T_n} \left[\log \mathbb{P}(\mathbf{u}_n^{\tau_{ns}} | \mathbf{y}_n^{\tau_{ns}}, \eta) + \log \mathbb{P}(\mathbf{y}_n^{\tau_{ns}} | \mathbf{z}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{n\,s-1}}, \mu) \right] \mathbb{I}\{s \in \mathcal{O}_n\}$$
$$+ \sum_{s=1}^{T_n} \log \mathbb{P}(\mathbf{u}_n^{\tau_{ns}} | \mathbf{z}_n^{\tau_{ns}}, \eta^{'})\, \mathbb{I}\{s \in \mathcal{O}_n^{'}\} + \sum_{s=1}^{T_n - 1} \left[\log \mathbb{P}(\mathbf{z}_n^{\tau_n(s)} | \mathbf{z}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}}, \rho, \delta) \right]$$

where $\mathbf{u}_n^{\tau_{n0}} = 0$ is assumed as the initial intervention before the first visit. By defining the variables $\chi_{n\,ik}^{\tau_n(s)}$ and $\psi_{n\,i}^{\tau_n(s)}$ as in the baseline model, the log probabilities can be written as:

$$\log \mathbb{P}(\mathbf{z}_n^{\tau_n(s)} | \mathbf{z}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}}, \rho, \delta) = \sum_{i=1}^{I} \sum_{\substack{k=1 \\ k \neq i}}^{I} \left(\chi_{n\,ik}^{\tau_n(s)} \log([\mathcal{Q}_{n(s)}]_{ik}) - \psi_{n\,i}^{\tau_n(s)} [\mathcal{Q}_{n(s)}]_{ik} \right)$$

$$\log \mathbb{P}(\mathbf{y}_n^{\tau_{ns}} | \mathbf{z}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{n\,s-1}}, \mu) = \sum_{i=1}^{I} \sum_{l=0}^{L-1} \left[\log \binom{J-1}{\mathbf{y}_n^{\tau_{ns}}} + \mathbf{y}_n^{\tau_{ns}} \log \mu_{il} + (J - 1 - \mathbf{y}_n^{\tau_{ns}}) \log(1 - \mu_{il}) \right]$$
$$\times\, \mathbb{I}\{\mathbf{z}_n^{\tau_{ns}} = i\}\, \mathbb{I}\{\mathbf{u}_n^{\tau_{n\,s-1}} = l\}$$

$$\log \mathbb{P}(\mathbf{u}_n^{\tau_{ns}} | \mathbf{y}_n^{\tau_{ns}}, \eta) = \sum_{j=0}^{J-1} \left[\log \binom{L-1}{\mathbf{u}_n^{\tau_{ns}}} + \mathbf{u}_n^{\tau_{ns}} \log \eta_j + (L - 1 - \mathbf{u}_n^{\tau_{ns}}) \log(1 - \eta_j) \right] \mathbb{I}\{\mathbf{y}_n^{\tau_{ns}} = j\}$$

$$\log \mathbb{P}(\mathbf{u}_n^{\tau_{ns}} | \mathbf{z}_n^{\tau_{ns}}, \eta^{'}) = \sum_{i=1}^{I} \left[\log \binom{L-1}{\mathbf{u}_n^{\tau_{ns}}} + \mathbf{u}_n^{\tau_{ns}} \log \eta_i^{'} + (L - 1 - \mathbf{u}_n^{\tau_{ns}}) \log(1 - \eta_i^{'}) \right] \mathbb{I}\{\mathbf{z}_n^{\tau_{ns}} = i\}$$

The complete-data log-likelihood for all the patients will then be:

$$\log\, \mathcal{L}(\Theta) = \sum_{n=1}^{N} \log\, \mathcal{L}_n(\Theta) = \sum_{n=1}^{N} \sum_{i=1}^{I} \log(\pi_i)\, \mathbb{I}\{\mathbf{z}_n^{\tau_{n1}} = i\}$$
$$+ \sum_{n=1}^{N} \sum_{s=1}^{T_n - 1} \sum_{i=1}^{I} \sum_{\substack{k=1 \\ k \neq i}}^{I} \left[\chi_{n\,ik}^{\tau_n(s)} \left(\log(\delta_{ik\,\mathbf{u}_n^{\tau_{ns}}}) + \rho_{ik\,\mathbf{u}_n^{\tau_{ns}}}^{'} w_n \right) - \psi_{n\,i}^{\tau_n(s)}\, \delta_{ik\,\mathbf{u}_n^{\tau_{ns}}}\, \exp(\rho_{ik\,\mathbf{u}_n^{\tau_{ns}}}^{'} w_n) \right]$$

$$+\sum_{n=1}^{N}\sum_{s=1}^{T_n}\sum_{i=1}^{I}\sum_{l=0}^{L-1}\left[\log\binom{J-1}{\mathbf{y}_n^{\tau_{ns}}}+\mathbf{y}_n^{\tau_{ns}}\log(\mu_{il})+(J-1-\mathbf{y}_n^{\tau_{ns}})\log(1-\mu_{il})\right]$$
$$\times\,\mathbb{I}\{\mathbf{z}_n^{\tau_{ns}}=i,\,\mathbf{u}_n^{\tau_n\,s-1}=l,\,s\in\mathcal{O}_n\}$$
$$+\sum_{n=1}^{N}\sum_{s=1}^{T_n}\sum_{j=0}^{J-1}\left[\log\binom{L-1}{\mathbf{u}_n^{\tau_{ns}}}+\mathbf{u}_n^{\tau_{ns}}\log(\eta_j)+(L-1-\mathbf{u}_n^{\tau_{ns}})\log(1-\eta_j)\right]\mathbb{I}\{\mathbf{y}_n^{\tau_{ns}}=j,\,s\in\mathcal{O}_n\}$$
$$+\sum_{n=1}^{N}\sum_{s=1}^{T_n}\sum_{i=1}^{I}\left[\log\binom{L-1}{\mathbf{u}_n^{\tau_{ns}}}+\mathbf{u}_n^{\tau_{ns}}\log(\eta_i^{'})+(L-1-\mathbf{u}_n^{\tau_{ns}})\log(1-\eta_i^{'})\right]\mathbb{I}\{\mathbf{z}_n^{\tau_{ns}}=i,\,s\in\mathcal{O}_n^{'}\}$$

Let's take expectation with respect to the posterior distribution of the unobserved variables:

$$\mathbb{E}[\log\,\mathcal{L}(\Theta)|\mathbf{u}_n^{\tau_n\,obs},\mathbf{y}_n^{\tau_n\,obs},\mathbf{z}_n^{\tau_n\,obs},\tilde{\Theta}]=\sum_{n=1}^{N}\sum_{i=1}^{I}\log(\pi_i)\,\mathbb{I}\{\mathbf{z}_n^{\tau_{n1}}=i,\,1\in\mathcal{O}_n^{'}\}$$
$$+\sum_{n=1}^{N}\sum_{i=1}^{I}\log(\pi_i)\,\mathbb{P}(\mathbf{z}_n^{\tau_{n1}}=i\,|\,\mathbf{u}_n^{\tau_n\,obs},\mathbf{y}_n^{\tau_n\,obs},\mathbf{z}_n^{\tau_n\,obs},\tilde{\Theta})\,\mathbb{I}\{1\in\mathcal{O}_n\}$$
$$+\sum_{n=1}^{N}\sum_{s=1}^{T_n-1}\sum_{i=1}^{I}\sum_{\substack{k=1\\k\neq i}}^{I}\sum_{l=0}^{L-1}\Big[\mathbb{E}\big[\chi_{n\,ik}^{\tau_n\,(s)}|\mathbf{u}_n^{\tau_n\,obs},\mathbf{y}_n^{\tau_n\,obs},\mathbf{z}_n^{\tau_n\,obs},\tilde{\Theta}\big]\big(\log(\delta_{ikl})+\rho_{ikl}^{'}w_n\big)$$
$$-\,\mathbb{E}[\psi_{n\,i}^{\tau_n\,(s)}|\mathbf{u}_n^{\tau_n\,obs},\mathbf{y}_n^{\tau_n\,obs},\mathbf{z}_n^{\tau_n\,obs},\tilde{\Theta}]\,\delta_{ikl}\,\exp(\rho_{ikl}^{'}w_n)\Big]\mathbb{I}\{\mathbf{u}_n^{\tau_{ns}}=l\}$$
$$+\sum_{n=1}^{N}\sum_{s=1}^{T_n}\sum_{i=1}^{I}\sum_{l=0}^{L-1}\left[\log\binom{J-1}{\mathbf{y}_n^{\tau_{ns}}}+\mathbf{y}_n^{\tau_{ns}}\log(\mu_{il})+(J-1-\mathbf{y}_n^{\tau_{ns}})\log(1-\mu_{il}))\right]$$
$$\times\,\mathbb{P}(\mathbf{z}_n^{\tau_{ns}}=i\,|\,\mathbf{u}_n^{\tau_n\,obs},\mathbf{y}_n^{\tau_n\,obs},\mathbf{z}_n^{\tau_n\,obs},\tilde{\Theta})\,\mathbb{I}\{\mathbf{u}_n^{\tau_n\,s-1}=l,\,s\in\mathcal{O}_n\}$$
$$+\sum_{n=1}^{N}\sum_{s=1}^{T_n}\sum_{j=0}^{J-1}\left[\log\binom{L-1}{\mathbf{u}_n^{\tau_{ns}}}+\mathbf{u}_n^{\tau_{ns}}\log(\eta_j)+(L-1-\mathbf{u}_n^{\tau_{ns}})\log(1-\eta_j)\right]\mathbb{I}\{\mathbf{y}_n^{\tau_{ns}}=j,\,s\in\mathcal{O}_n\}$$
$$+\sum_{n=1}^{N}\sum_{s=1}^{T_n}\sum_{j=0}^{J-1}\left[\log\binom{L-1}{\mathbf{u}_n^{\tau_{ns}}}+\mathbf{u}_n^{\tau_{ns}}\log(\eta_j^{'})+(L-1-\mathbf{u}_n^{\tau_{ns}})\log(1-\eta_j^{'})\right]\mathbb{I}\{\mathbf{z}_n^{\tau_{ns}}=i,\,s\in\mathcal{O}_n^{'}\}$$

Assuming that we've calculated the above posterior probability and conditional expectations, the updated parameters in the E-step can be calculated by maximizing the expected complete-data log-likelihood with respect to Θ. Define $\gamma_n^s(i)=\mathbb{E}(\mathbb{I}\{\mathbf{z}_n^{\tau_{ns}}=i\}\,|\,\mathbf{u}_n^{\tau_n\,obs},\mathbf{y}_n^{\tau_n\,obs},\mathbf{z}_n^{\tau_n\,obs},\tilde{\Theta})$ (for $1\leq s\leq T_n$) and $\nu_n^s(i,k)=\mathbb{E}(\mathbb{I}\{\mathbf{z}_n^{\tau_{ns}}=i,\mathbf{z}_n^{\tau_n\,s+1}=k\}\,|\,\mathbf{u}_n^{\tau_n\,obs},\mathbf{y}_n^{\tau_n\,obs},\mathbf{z}_n^{\tau_n\,obs},\tilde{\Theta})$ (for

$1 \leq s \leq T_n - 1$). Notice that γ simplifies to an indicator function or a probability, respectively, if z has been observed or has not been observed. Similarly, ν can be simplified based on whether z has been observed at time periods s and $s+1$. In mathematical terms:

$$
\begin{aligned}
\gamma_n^s(i) &= \mathbb{I}\{\mathbf{z}_n^{\tau_{ns}} = i,\ s \in \mathcal{O}_n^{'}\} + \mathbb{P}(\mathbf{z}_n^{\tau_{ns}} = i \,|\, \mathbf{u}_n^{\tau_{n\,obs}}, \mathbf{y}_n^{\tau_{n\,obs}}, \mathbf{z}_n^{\tau_{n\,obs}}, \tilde{\Theta})\, \mathbb{I}\{s \in \mathcal{O}_n\} \\
\nu_n^s(i,k) &= \mathbb{I}\{\mathbf{z}_n^{\tau_{ns}} = i, \mathbf{z}_n^{\tau_{n\,s+1}} = k, s \in \mathcal{O}_n^{'}, s+1 \in \mathcal{O}_n^{'}\} \\
&\quad + \mathbb{P}(\mathbf{z}_n^{\tau_{n\,s+1}} = k \,|\, \mathbf{u}_n^{\tau_{n\,obs}}, \mathbf{y}_n^{\tau_{n\,obs}}, \mathbf{z}_n^{\tau_{n\,obs}}, \tilde{\Theta})\, \mathbb{I}\{\mathbf{z}_n^{\tau_{ns}} = i, s \in \mathcal{O}_n^{'}, s+1 \in \mathcal{O}_n\} \\
&\quad + \mathbb{P}(\mathbf{z}_n^{\tau_{ns}} = i \,|\, \mathbf{u}_n^{\tau_{n\,obs}}, \mathbf{y}_n^{\tau_{n\,obs}}, \mathbf{z}_n^{\tau_{n\,obs}}, \tilde{\Theta})\, \mathbb{I}\{\mathbf{z}_n^{\tau_{n\,s+1}} = k, s \in \mathcal{O}_n, s+1 \in \mathcal{O}_n^{'}\} \\
&\quad + \mathbb{P}(\mathbf{z}_n^{\tau_{ns}} = i, \mathbf{z}_n^{\tau_{n\,s+1}} = k \,|\, \mathbf{u}_n^{\tau_{n\,obs}}, \mathbf{y}_n^{\tau_{n\,obs}}, \mathbf{z}_n^{\tau_{n\,obs}}, \tilde{\Theta})\, \mathbb{I}\{s \in \mathcal{O}_n, s+1 \in \mathcal{O}_n\}
\end{aligned}
$$

For π, the optimal value is determined as $\hat{\pi}_i = \frac{\sum_{n=1}^{N} \gamma_n^1(i)}{\sum_{i'=1}^{I} \sum_{n=1}^{N} \gamma_n^1(i')}$ (for $1 \leq i \leq I$). Similarly, the optimal values for μ, η, and $\eta^{'}$ will be:

$$
\begin{aligned}
\hat{\mu}_{il} &= \frac{\sum_{n=1}^{N} \sum_{s=1}^{T_n} \mathbf{y}_n^{\tau_{ns}}\, \gamma_n^s(i)\, \mathbb{I}\{\mathbf{u}_n^{\tau_{n\,s-1}} = l,\ s \in \mathcal{O}_n\}}{(J-1) \sum_{n=1}^{N} \sum_{s=1}^{T_n} \gamma_n^s(i)\, \mathbb{I}\{\mathbf{u}_n^{\tau_{n\,s-1}} = l,\ s \in \mathcal{O}_n\}}; \quad (1 \leq i \leq I,\ 0 \leq l \leq L-1) \\
\hat{\eta}_j &= \frac{\sum_{n=1}^{N} \sum_{s=1}^{T_n} \mathbf{u}_n^{\tau_{ns}} \mathbb{I}\{\mathbf{y}_n^{\tau_{ns}} = j,\ s \in \mathcal{O}_n\}}{(L-1) \sum_{n=1}^{N} \sum_{s=1}^{T_n} \mathbb{I}\{\mathbf{y}_n^{\tau_{ns}} = j,\ s \in \mathcal{O}_n\}}; \quad (0 \leq j \leq J-1) \\
\hat{\eta}_i^{'} &= \frac{\sum_{n=1}^{N} \sum_{s=1}^{T_n} \mathbf{u}_n^{\tau_{ns}} \mathbb{I}\{\mathbf{z}_n^{\tau_{ns}} = i,\ s \in \mathcal{O}_n^{'}\}}{(L-1) \sum_{n=1}^{N} \sum_{s=1}^{T_n} \mathbb{I}\{\mathbf{z}_n^{\tau_{ns}} = i,\ s \in \mathcal{O}_n^{'}\}}; \quad (1 \leq i \leq I)
\end{aligned}
$$

For updating δ and ρ, equations 3.6, 3.7, and 3.8 can be used in their original form.

Therefore, in the E-step we need to calculate γ and ν as well as the two conditional expectations. We further decompose the conditional expectations as:

$$
\begin{aligned}
&\mathbb{E}[\chi_{n\,ik}^{\tau_n(s)} | \mathbf{u}_n^{\tau_{n\,obs}}, \mathbf{y}_n^{\tau_{n\,obs}}, \mathbf{z}_n^{\tau_{n\,obs}}, \tilde{\Theta}] = \\
&\qquad \sum_{i'=1}^{I} \sum_{k'=1}^{I} \mathbb{E}[\chi_{n\,ik}^{\tau_n(s)} | \mathbf{z}_n^{\tau_{ns}} = i^{'}, \mathbf{z}_n^{\tau_{n\,s+1}} = k^{'}, \mathbf{u}_n^{\tau_{n\,obs}}, \mathbf{y}_n^{\tau_{n\,obs}}, \mathbf{z}_n^{\tau_{n\,obs}}, \tilde{\Theta}]\, \nu_n^s(i', k') \\
&\mathbb{E}[\psi_{n\,i}^{\tau_n(s)} | \mathbf{u}_n^{\tau_{n\,obs}}, \mathbf{y}_n^{\tau_{n\,obs}}, \mathbf{z}_n^{\tau_{n\,obs}}, \tilde{\Theta}] =
\end{aligned}
$$

$$\sum_{i'=1}^{I}\sum_{k'=1}^{I}\mathbb{E}[\psi_{n\,i}^{\tau_{n}(s)}|\mathbf{z}_{n}^{\tau_{ns}}=i^{'},\mathbf{z}_{n}^{\tau_{n}\,s+1}=k^{'},\mathbf{u}_{n}^{\tau_{n}\,obs},\mathbf{y}_{n}^{\tau_{n}\,obs},\mathbf{z}_{n}^{\tau_{n}\,obs},\tilde{\Theta}]\,\nu_{n}^{s}(i',k')$$

The end-state conditioned expectations can be calculated similar to the previous section. Consequently, the main remaining challenge would be to calculate γ and ν, which are respectively equivalent to calculating $\mathbb{P}(\mathbf{z}_{n}^{\tau_{ns}}=i\,|\,\mathbf{u}_{n}^{\tau_{n}\,obs},\mathbf{y}_{n}^{\tau_{n}\,obs},\mathbf{z}_{n}^{\tau_{n}\,obs},\tilde{\Theta})$ for $s\in\mathcal{O}_{n}$ and $\mathbb{P}(\mathbf{z}_{n}^{\tau_{ns}}=i,\mathbf{z}_{n}^{\tau_{n}\,s+1}=k\,|\,\mathbf{u}_{n}^{\tau_{n}\,obs},\mathbf{y}_{n}^{\tau_{n}\,obs},\mathbf{z}_{n}^{\tau_{n}\,obs},\tilde{\Theta})$ for $s,s+1\in\mathcal{O}_{n}$. In the following, we develop a forward-backward algorithm for efficiently calculating these probabilities.

3.6.3 The Forward-Backward Algorithm

Similar to the previous case, construct a time-inhomogeneous DT-HMM from the CT-HMM by defining the transition probability function of the DT-HMM to be equal to $[\tilde{\mathcal{P}}_{n\,(s)}]_{ik}=\mathbb{P}(\tilde{\mathbf{z}}_{n}^{s+1}=k|\tilde{\mathbf{z}}_{n}^{s}=i,\mathbf{u}_{n}^{\tau_{ns}},\tilde{\Theta})=\left[e^{(\tau_{n\,s+1}-\tau_{ns})\tilde{\mathcal{Q}}_{n\,(s)}}\right]_{ik}$ where $\tilde{\mathbf{z}}$ is the discrete-time latent variable of our DT-HMM and $\tilde{\mathcal{Q}}_{n\,(s)}$ is the generator matrix, associated with interval $(\tau_{ns},\tau_{n\,s+1}]$. For a given time period s, denote by $\mathbf{z}_{n}^{\tau_{n}\,(s)+}$ the set of underlying health condition variables after visit s that correspond to direct observation through accurate examination. In other words, $\mathbf{z}_{n}^{\tau_{n}\,(s)+}$ includes all $\mathbf{z}_{n}^{\tau_{ns'}}$ variables such that $s^{'}>s$ and $s^{'}\in\mathcal{O}_{n}^{'}$. Moreover, let $\mathbf{y}_{n}^{\tau_{n}\,(s)+}$ denote physician observation variables after visit s, i.e. $\mathbf{y}_{n}^{\tau_{n}\,s+}$ includes all $\mathbf{y}_{n}^{\tau_{n}\,s'}$ for which $s'>s$ and $s\in\mathcal{O}_{n}$. Finally, let $\mathbf{u}_{n}^{\tau_{n}\,(s)+}=\mathbf{u}_{n}^{\tau_{n}\,s+1:T_{n}}$ be the set of interventions after visit s. From a conceptual point of view, this group consists of all the variables after visit s for which the values are known. Similarly, define the second group $\mathbf{z}_{n}^{\tau_{n}\,(s)-}$, $\mathbf{y}_{n}^{\tau_{n}\,(s)-}$, and $\mathbf{u}_{n}^{\tau_{n}\,(s)-}$ as the corresponding variables to the previous group for at or before time period s. Notice that the union of these two sets covers all the variables for which the values are known. The forward probability associated with visit s will then be $\alpha_{n}^{s}(i)=\mathbb{P}(\tilde{\mathbf{z}}_{n}^{s}=i,\mathbf{u}_{n}^{\tau_{n}\,(s)-},\mathbf{y}_{n}^{\tau_{n}\,(s)-},\mathbf{z}_{n}^{\tau_{n}\,(s)-}|\tilde{\Theta})$. We define $\alpha_{n}^{s}(i)$ only for $i=\tilde{\mathbf{z}}_{n}^{s}$ when the underlying condition at time point s has been directly observed via accurate examination (i.e. $s\in\mathcal{O}_{n}^{'}$). The backward probability is also defined as $\beta_{n}^{s}(i)=\mathbb{P}(\mathbf{u}_{n}^{\tau_{n}\,(s)+},\mathbf{y}_{n}^{\tau_{n}\,(s)+},\mathbf{z}_{n}^{\tau_{n}\,(s)+}|\tilde{\mathbf{z}}_{n}^{s}=i,\mathbf{u}_{n}^{\tau_{ns}},\tilde{\Theta})$ in a similar way.

In the following, we'll present a set of dynamic programming equations that allow for efficient calculation of α and β. Depending on whether $s \in \mathcal{O}_n$ and $s-1 \in \mathcal{O}_n$, there are four cases for calculation of α. More specifically:

$$
\begin{aligned}
s-1 \in \mathcal{O}_n, s \in \mathcal{O}_n : \alpha_n^s(i) &= \sum_{k=1}^{I} \mathbb{P}(\tilde{\mathbf{z}}_n^s = i, \tilde{\mathbf{z}}_n^{s-1} = k, \mathbf{u}_n^{\tau_n(s)-}, \mathbf{y}_n^{\tau_n(s)-}, \mathbf{z}_n^{\tau_n(s)-} | \tilde{\Theta}) \\
&= \sum_{k=1}^{I} \mathbb{P}(\mathbf{u}_n^{\tau_{ns}} | \mathbf{y}_n^{\tau_{ns}}, \tilde{\Theta}) \mathbb{P}(\mathbf{y}_n^{\tau_{ns}} | \tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_n s-1}, \tilde{\Theta}) \times \\
&\quad \mathbb{P}(\tilde{\mathbf{z}}_n^s = i | \tilde{\mathbf{z}}_n^{s-1} = k, \mathbf{u}_n^{\tau_n s-1}, \tilde{\Theta}) \, \mathbb{P}(\tilde{\mathbf{z}}_n^{s-1} = k, \mathbf{u}_n^{\tau_n(s-1)-}, \mathbf{y}_n^{\tau_n(s-1)-}, \mathbf{z}_n^{\tau_n(s-1)-} | \tilde{\Theta}) \\
&= \sum_{k=1}^{I} \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}}) \, \tilde{\mathcal{E}}(i, \mathbf{u}_n^{\tau_n s-1}, \mathbf{y}_n^{\tau_{ns}}) \left[\tilde{\mathcal{P}}_{n(s-1)}\right]_{ki} \alpha_n^{s-1}(k) \qquad (3.22)
\end{aligned}
$$

$$
s-1 \in \mathcal{O}_n^{'}, s \in \mathcal{O}_n : \alpha_n^s(i) = \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}}) \, \tilde{\mathcal{E}}(i, \mathbf{u}_n^{\tau_n s-1}, \mathbf{y}_n^{\tau_{ns}}) \left[\tilde{\mathcal{P}}_{n(s-1)}\right]_{\tilde{\mathbf{z}}_n^{s-1} i} \alpha_n^{s-1}(\tilde{\mathbf{z}}_n^{s-1}) \qquad (3.23)
$$

$$
s-1 \in \mathcal{O}_n, s \in \mathcal{O}_n^{'} : \alpha_n^s(\tilde{\mathbf{z}}_n^s) = \sum_{k=1}^{I} \tilde{\mathcal{G}}^{'}(\tilde{\mathbf{z}}_n^s, \mathbf{u}_n^{\tau_{ns}}) \left[\tilde{\mathcal{P}}_{n(s-1)}\right]_{k \, \tilde{\mathbf{z}}_n^s} \alpha_n^{s-1}(k) \qquad (3.24)
$$

$$
s-1 \in \mathcal{O}_n^{'}, s \in \mathcal{O}_n^{'} : \alpha_n^s(\tilde{\mathbf{z}}_n^s) = \tilde{\mathcal{G}}^{'}(\tilde{\mathbf{z}}_n^s, \mathbf{u}_n^{\tau_{ns}}) \left[\tilde{\mathcal{P}}_{n(s-1)}\right]_{\tilde{\mathbf{z}}_n^{s-1} \, \tilde{\mathbf{z}}_n^s} \alpha_n^{s-1}(\tilde{\mathbf{z}}_n^{s-1}) \qquad (3.25)
$$

For the calculation of β, we have two cases based on $s+1 \in \mathcal{O}_n$:

$$
\begin{aligned}
s+1 \in \mathcal{O}_n : \beta_n^s(i) &= \sum_{k=1}^{I} \mathbb{P}(\tilde{\mathbf{z}}_n^{s+1} = k, \mathbf{u}_n^{\tau_n(s)+}, \mathbf{y}_n^{\tau_n(s)+}, \mathbf{z}_n^{\tau_n(s)+} | \tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_{ns}}, \tilde{\Theta}) \\
&= \sum_{k=1}^{I} \mathbb{P}(\mathbf{u}_n^{\tau_n(s+1)+}, \mathbf{y}_n^{\tau_n(s+1)+}, \mathbf{z}_n^{\tau_n(s+1)+} | \tilde{\mathbf{z}}_n^{s+1} = k, \mathbf{u}_n^{\tau_n s+1}, \tilde{\Theta}) \, \mathbb{P}(\mathbf{u}_n^{\tau_n s+1} | \mathbf{y}_n^{\tau_n s+1}, \tilde{\Theta}) \\
&\quad \mathbb{P}(\mathbf{y}_n^{\tau_n s+1} | \tilde{\mathbf{z}}_n^{s+1} = k, \mathbf{u}_n^{\tau_n s}, \tilde{\Theta}) \, \mathbb{P}(\tilde{\mathbf{z}}_n^{s+1} = k | \tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_{ns}}, \tilde{\Theta}) \\
&= \sum_{k=1}^{I} \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_n s+1}, \mathbf{u}_n^{\tau_n s+1}) \, \tilde{\mathcal{E}}(k, \mathbf{u}_n^{\tau_n s}, \mathbf{y}_n^{\tau_n s+1}) \left[\tilde{\mathcal{P}}_{n(s)}\right]_{ik} \beta_n^{s+1}(k) \qquad (3.26)
\end{aligned}
$$

$$
s+1 \in \mathcal{O}_n^{'} : \beta_n^s(i) = \tilde{\mathcal{G}}^{'}(\tilde{\mathbf{z}}_n^{s+1}, \mathbf{u}_n^{\tau_n s+1}) \left[\tilde{\mathcal{P}}_{n(s)}\right]_{i \, \tilde{\mathbf{z}}_n^{s+1}} \beta_n^{s+1}(\tilde{\mathbf{z}}_n^{s+1}) \qquad (3.27)
$$

Notice that $\beta_n^s(i)$ is defined only for $i = \tilde{\mathbf{z}}_n^s$ if $s \in \mathcal{O}_n^{'}$. The boundary conditions will also be:

$$1 \in \mathcal{O}_n : \alpha_n^1(i) = \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_{n1}}, \mathbf{u}_n^{\tau_{n1}})\, \tilde{\mathcal{E}}(i, \mathbf{u}_n^{\tau_n 0}, \mathbf{y}_n^{\tau_{n1}}) \tilde{\pi}_i \tag{3.28}$$

$$1 \in \mathcal{O}_n^{'} : \alpha_n^1(\tilde{\mathbf{z}}_n^1) = \tilde{\mathcal{G}}^{'}(\tilde{\mathbf{z}}_n^1, \mathbf{u}_n^{\tau_{n1}})\, \tilde{\pi}_{\tilde{\mathbf{z}}_n^1} \tag{3.29}$$

$$T_n \in \mathcal{O}_n : \beta_n^{T_n-1}(i) = \sum_{k=1}^{I} \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_n T_n}, \mathbf{u}_n^{\tau_n T_n})\, \tilde{\mathcal{E}}(k, \mathbf{u}_n^{\tau_n T_n - 1}, \mathbf{y}_n^{\tau_n T_n}) \left[\tilde{\mathcal{P}}_{n\,(T_n-1)}\right]_{ik} \tag{3.30}$$

$$T_n \in \mathcal{O}_n^{'} : \beta_n^{T_n-1}(i) = \tilde{\mathcal{G}}(\tilde{\mathbf{z}}_n^{T_n}, \mathbf{u}_n^{\tau_n T_n}) \left[\tilde{\mathcal{P}}_{n\,(T_n-1)}\right]_{i\,\tilde{\mathbf{z}}_n^{T_n}} \tag{3.31}$$

where $\beta_n^{T_n-1}(i)$ is defined only for $i = \tilde{\mathbf{z}}_n^{T_n-1}$ if $T_n - 1 \in \mathcal{O}_n^{'}$. Now, we can calculate the posterior probabilities that we are interested in as:

$$\mathbb{P}(\mathbf{z}_n^{\tau_{ns}} = i \,|\, \mathbf{u}_n^{\tau_n\, obs}, \mathbf{y}_n^{\tau_n\, obs}, \mathbf{z}_n^{\tau_n\, obs}, \tilde{\Theta}) = \frac{\alpha_n^s(i)\beta_n^s(i)}{\sum_{i'=1}^{I} \alpha_n^s(i')\beta_n^s(i')} \tag{3.32}$$

$$\mathbb{P}(\mathbf{z}_n^{\tau_{ns}} = i, \mathbf{z}_n^{\tau_n\, s+1} = k \,|\, \mathbf{u}_n^{\tau_n\, obs}, \mathbf{y}_n^{\tau_n\, obs}, \mathbf{z}_n^{\tau_n\, obs}, \tilde{\Theta}) =$$
$$\frac{\tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_n\, s+1}, \mathbf{u}_n^{\tau_n\, s+1})\, \tilde{\mathcal{E}}(k, \mathbf{u}_n^{\tau_n\, s}, \mathbf{y}_n^{\tau_n\, s+1}) \left[\tilde{\mathcal{P}}_{n\,(s)}\right]_{ik}\, \alpha_n^s(i)\, \beta_n^{s+1}(k)}{\sum_{i'=1}^{I}\sum_{k'=1}^{I} \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_n\, s+1}, \mathbf{u}_n^{\tau_n\, s+1})\, \tilde{\mathcal{E}}(k', \mathbf{u}_n^{\tau_n\, s}, \mathbf{y}_n^{\tau_n\, s+1}) \left[\tilde{\mathcal{P}}_{n\,(s)}\right]_{i'\,k'}\, \alpha_n^s(i')\, \beta_n^{s+1}(k')} \tag{3.33}$$

The boundary cases need to be handled separately, similar to the baseline model.

3.6.4 Numerical Experimentation

We generate synthetic data to validate our algorithm as in Section 3.4.5[2]. Specifically, we assume that physicians order accurate examination with a certain probability whenever the observation in a regular visit indicates the patient's health condition suggests that the disease has progressed to high-risk states. Consequently, a snapshot of the underlying health condition will be available in the next visit. Similarly, the physician decides whether to repeat the accurate examination for the next period or go back to the regular-visits regimen based on the observed underlying health condition. We model the probability according to which

[2]Refer to directory CTHMM-Ext2 in *https://github.com/saeedghodsi93/Disease_Progression_Modeling_HMM*

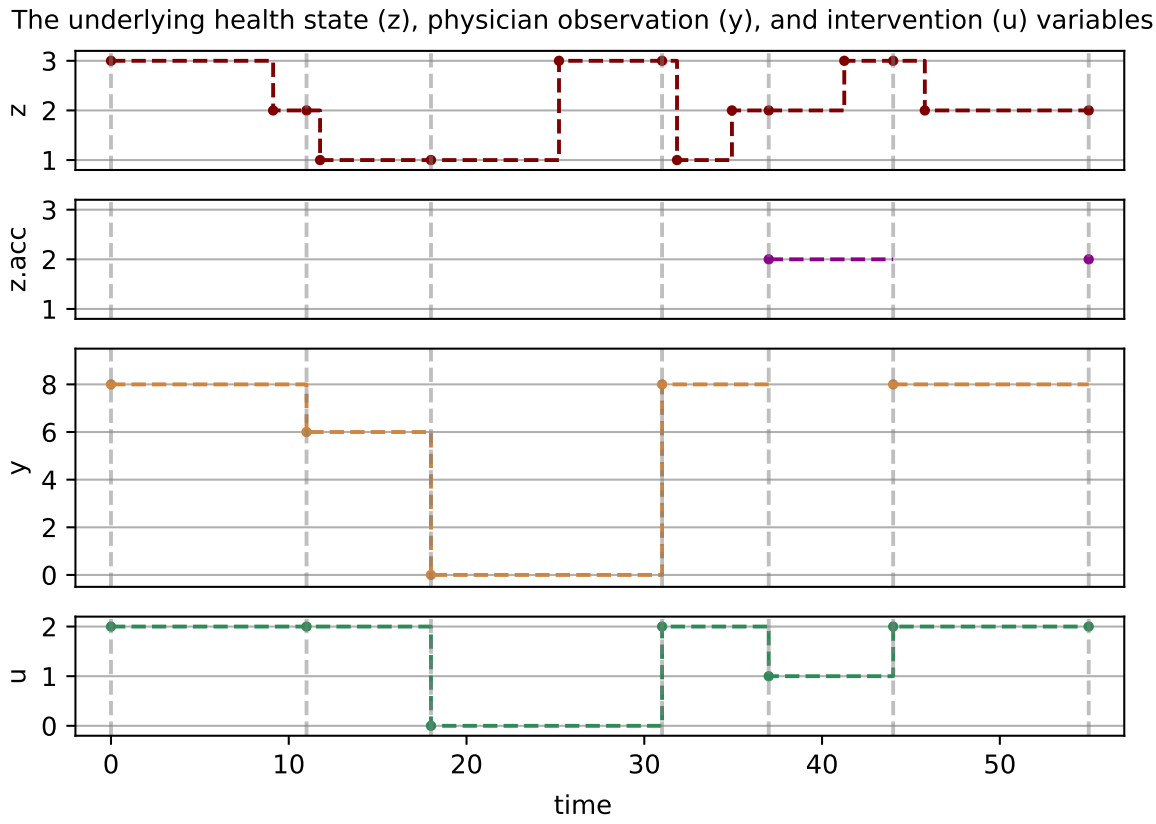

Figure 3.5: An example of the generated synthetic data. In periods that correspond to accurate examination, the value of z is observed instead of y.

the accurate-examination decision at a given time period is made as a logistic function of the observation that was made in either the regular or accurate examination during the previous visit. In mathematical terms, $\mathbb{P}(s \in \mathcal{O}_n^{'}) = \frac{exp(\Upsilon)}{1+exp(\Upsilon)}$, where Υ is an increasing linear function of $\mathbf{y}_n^{\tau_n s-1}$ or $\mathbf{z}_n^{\tau_n s-1}$, depending on the case. The rest of the data generation procedure remains almost the same. The only difference with the baseline case is that this time we generate $\mathbf{y}_n^{\tau_{ns}}$ based on a binomial distribution with parameter μ_{il} for $i = \mathbf{z}_n^{\tau_{ns}}$ and $l = \mathbf{u}_n^{\tau_n s-1}$. We assume $\mathbf{u}_n^{\tau_{n0}} = 0$, which means that the lowest level of intervention (possibly no intervention) was being administered until the first visit time.

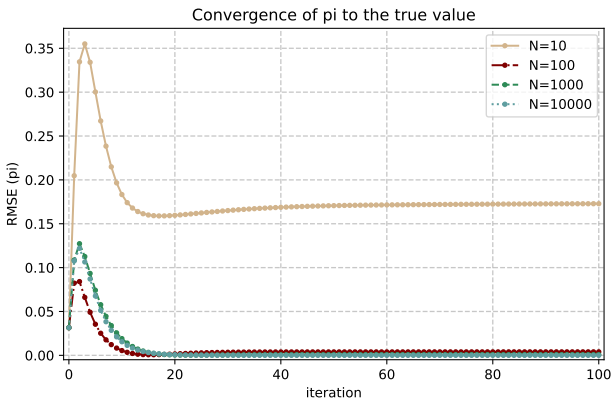

(a) Convergence of $\hat{\pi}$

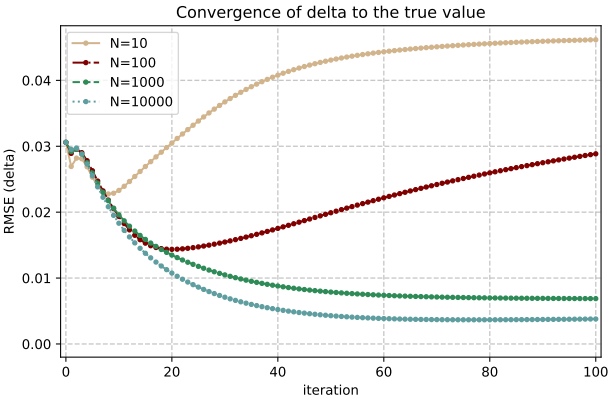

(b) Convergence of $\hat{\delta}$

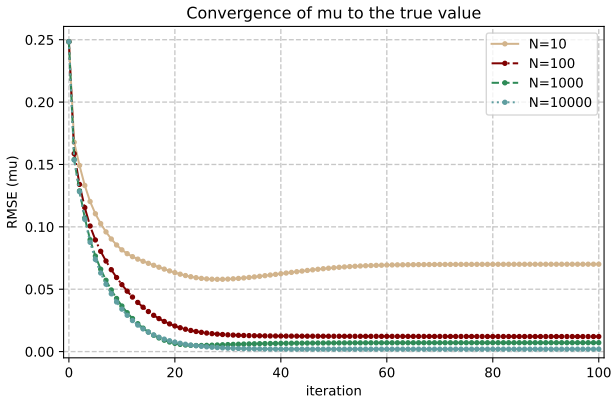

(c) Convergence of $\hat{\mu}$

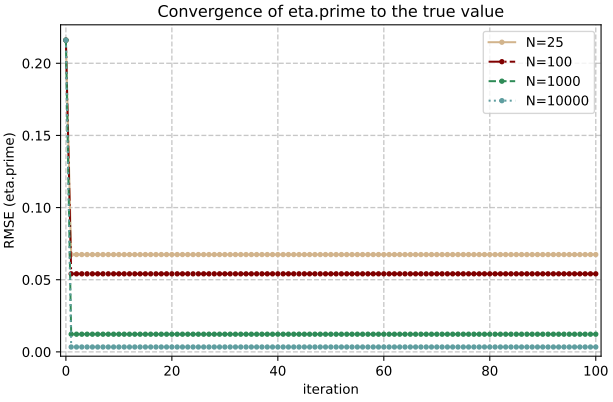

(d) Convergence of $\hat{\eta}^{'}$

Figure 3.6: Convergence of the EM estimated parameters to their true values for different sample sizes (the $\hat{\eta}$ plot is eliminated since it is constant after the first iteration). As the results indicate, convergence is faster than the baseline model (i.e. Figure 3.3).

Figure 3.5 illustrates an example of the generated synthetic data. We present the convergence plots for variables $\hat{\pi}$, $\hat{\delta}$, $\hat{\mu}$, and $\hat{\eta}'$ in Figure 3.6. One can argue that the convergence rates should be faster than the baseline model mainly due to the fact that being able to occasionally observe the underlying health condition provides the algorithm with more accurate information. However, having an additional link from $\mathbf{u}_n^{\tau_n s-1}$ to $\mathbf{y}_n^{\tau_n s}$ makes the structure of the model more complex and reduces the convergence rates as a result. We observe that with the same number of samples, the extended model converges slightly faster and the ultimate errors are smaller than the baseline model. Hence, the effect of having occasional access to the latent variable dominates the other effect. Indeed, the rate at which the latent variable is observed plays a key role in the algorithm's convergence rate.

3.7 Individual-Level Intervention Planning

In this section, we consider a patient-level planning scenario in which the history of the patient is used for estimating the posterior distribution of their underlying health state at the current time. The goal would then be to choose an intervention that minimizes the total expected disutility over the planning horizon. There are two main factors that contribute to the disutility: 1- the discomfort associated with being in a certain underlying health state, 2- the risk associated with receiving each intervention option. For example, the patient suffers from a certain level of pain while being in each stage of cancer. On the other hand, an aggressive intervention such as removing cancer as well as some healthy tissues around it via surgery puts the patient at the risk of side effects. Hence, we define the total disutility as the addition of two terms that each corresponds to one of the aforementioned components. Designing an appropriate and timely intervention plan requires keeping track of the changes in the underlying health state over time.

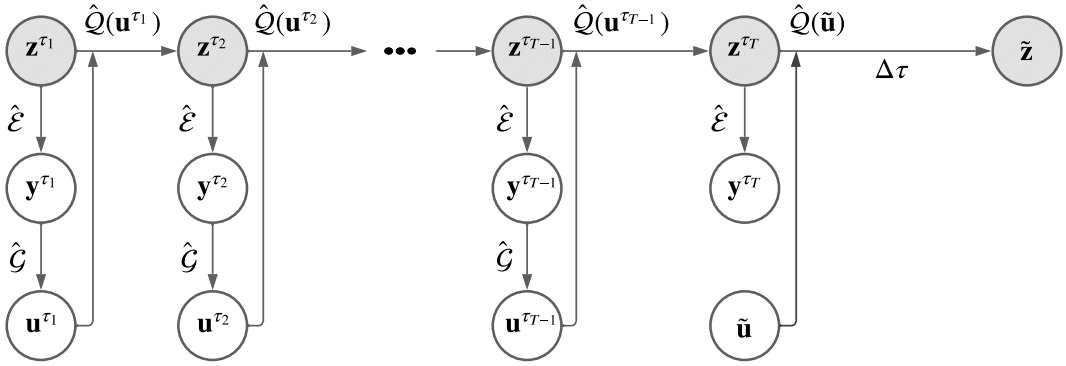

Figure 3.7: Given the patient's history (i.e. $\mathbf{y}^{\tau_{1:T}}$, $\mathbf{u}^{\tau_{1:T-1}}$) and the model parameters, the posterior distribution of the underlying health state $\mathbf{z}^{\tau_T}$ can be evaluated. The intervention $\mathbf{u}$, which affects the generator matrix $\hat{\mathcal{Q}}(\mathbf{u})$ can then be chosen in a way that the expected disutility is minimized. The disutility Ω is modeled as a function of the patient's underlying health state $\tilde{\mathbf{z}}$ at the end of the planning period (i.e. at time $\tau_T + \Delta\tau$).

3.7.1 Single-Period Planning

Suppose we have estimated the parameter set $\hat{\Theta} = (\hat{\pi}, \hat{\rho}, \hat{\delta}, \hat{\mu}, \hat{\eta})$ using a training dataset. Consider a new patient for which the history of physician observations $\mathbf{y}^{\tau_{1:T}}$ as well as the prescribed interventions $\mathbf{u}^{\tau_{1:T-1}}$ is available up to the current time period. Define $\Omega^{\mathbf{z}}$ as a measure of disutility that depends on the patient's underlying unobserved health state. Moreover, let $\Omega^{\mathbf{u}}$ indicate the disutility associated with the intervention riskiness. We define the total disutility as $\Omega = \Omega^{\mathbf{z}} + \Omega^{\mathbf{u}}$. In this section, we study the simplest version of the intervention planning problem, that is a single-period optimization problem. Specifically, we assume to have a planning period of fixed length $\Delta\tau$ and try to minimize the expected disutility at the end of the planning period. In mathematical terms, we would like to choose an intervention $\mathbf{u}$ at time period T that minimizes $\mathbb{E}[\Omega]$ during period $[\tau_T, \tau_T + \Delta\tau)$ (refer to Figure 3.7). Thus, we develop an optimization problem that minimizes the expected disutility, given the patient's history.

Let w indicates the covariate vector for the patient. The single-period intervention planning problem aims to find the optimal intervention at time period T, denoted by $\tilde{\mathbf{u}}$, given

w, $\mathbf{y}^{\tau_{1:T}}$, and $\mathbf{u}^{\tau_{1:T-1}}$. Therefore, we essentially need to evaluate the conditional expected disutility $\mathbb{E}[\Omega|\tilde{\mathbf{u}}, \mathbf{y}^{\tau_{1:T}}, \mathbf{u}^{\tau_{1:T-1}}, w, \hat{\Theta}]$ during period $[\tau_T,\ \tau_T + \Delta\tau)$, which can be broken to the sum of $\Omega^{\mathbf{u}}$ and $\mathbb{E}[\Omega^{\mathbf{z}}|\tilde{\mathbf{u}}, \mathbf{y}^{\tau_{1:T}}, \mathbf{u}^{\tau_{1:T-1}}, w, \hat{\Theta}]$ at time $\tau_T + \Delta\tau$. By applying the law of iterated expectations we can simplify the second term of the objective as $\mathbb{E}[\Omega^{\mathbf{z}}|\tilde{\mathbf{u}}, \mathbf{y}^{\tau_{1:T}}, \mathbf{u}^{\tau_{1:T-1}}, w, \hat{\Theta}] = \mathbb{E}\big[\mathbb{E}[\Omega^{\mathbf{z}}|\mathbf{z}^{\tau_T}] \,|\, \tilde{\mathbf{u}}, \mathbf{y}^{\tau_{1:T}}, \mathbf{u}^{\tau_{1:T-1}}, w, \hat{\Theta}\big]$, where the outer expectation is calculated with respect to the posterior probability of the underlying health state given the historical data, which is proportional to $\mathbb{P}(\mathbf{z}^{\tau_T}, \mathbf{y}^{\tau_{1:T}}, \mathbf{u}^{\tau_{1:T-1}} \,|\, w, \hat{\Theta})$. This joint probability can be factorized as:

$$
\begin{aligned}
&\mathbb{P}(\mathbf{z}^{\tau_T} = i, \mathbf{y}^{\tau_{1:T}}, \mathbf{u}^{\tau_{1:T-1}} \,|\, w, \hat{\Theta}) = \sum_{k=1}^{I} \mathbb{P}(\mathbf{z}^{\tau_T} = i, \mathbf{z}^{\tau_{T-1}} = k, \mathbf{y}^{\tau_{1:T}}, \mathbf{u}^{\tau_{1:T-1}} \,|\, w, \hat{\Theta}) \\
&= \sum_{k=1}^{I} \mathbb{P}(\mathbf{y}^{\tau_T} \,|\, \mathbf{z}^{\tau_T} = i, \hat{\Theta})\, \mathbb{P}(\mathbf{z}^{\tau_T} = i \,|\, \mathbf{z}^{\tau_{T-1}} = k, \mathbf{u}^{\tau_{T-1}}, w, \hat{\Theta})\, \mathbb{P}(\mathbf{z}^{\tau_{T-1}} = k, \mathbf{y}^{\tau_{1:T-1}}, \mathbf{u}^{\tau_{1:T-1}} \,|\, w, \hat{\Theta}) \\
&= \sum_{k=1}^{I} \hat{\mathcal{E}}(i, \mathbf{y}^{\tau_T}) \left[\hat{\mathcal{P}}_{(T-1)}\right]_{ki} \alpha^{T-1}(k)
\end{aligned}
$$

where $\left[\hat{\mathcal{P}}_{(T-1)}\right]_{ki}$ indicates the transition probability of the corresponding time-inhomogeneous DT-HMM in the time period $[\tau_{T-1}, \tau_T)$, and $\alpha^{T-1}(k)$ is the forward probability at time $T-1$, which can be efficiently calculated via Equations 3.9 and 3.11. The posterior probability of $\mathbf{z}^{\tau_T}$ would then be:

$$
\mathbb{P}(\mathbf{z}^{\tau_T} = i \,|\, \mathbf{y}^{\tau_{1:T}}, \mathbf{u}^{\tau_{1:T-1}}, w, \hat{\Theta}) = \frac{\sum_{k=1}^{I} \hat{\mathcal{E}}(i, \mathbf{y}^{\tau_T}) \left[\hat{\mathcal{P}}_{(T-1)}\right]_{ki} \alpha^{T-1}(k)}{\sum_{i'=1}^{I} \sum_{k=1}^{I} \hat{\mathcal{E}}(i', \mathbf{y}^{\tau_T}) \left[\hat{\mathcal{P}}_{(T-1)}\right]_{ki'} \alpha^{T-1}(k)} \tag{3.34}
$$

Hence, we can represent the aforementioned expected disutility as:

$$
\begin{aligned}
&\mathbb{E}[\Omega^{\mathbf{z}}|\tilde{\mathbf{u}}, \mathbf{y}^{\tau_{1:T}}, \mathbf{u}^{\tau_{1:T-1}}, w, \hat{\Theta}] = \sum_{i=1}^{I} \mathbb{E}[\Omega^{\mathbf{z}}|\mathbf{z}^{\tau_T} = i, \tilde{\mathbf{u}}, w, \hat{\Theta}]\, \mathbb{P}(\mathbf{z}^{\tau_T} = i \,|\, \mathbf{y}^{\tau_{1:T}}, \mathbf{u}^{\tau_{1:T-1}}, w, \hat{\Theta}) \\
&= \frac{\sum_{i=1}^{I} \sum_{k=1}^{I} \mathbb{E}[\Omega^{\mathbf{z}}|\mathbf{z}^{\tau_T} = i, \tilde{\mathbf{u}}, w, \hat{\Theta}]\, \hat{\mathcal{E}}(i, \mathbf{y}^{\tau_T}) \left[\hat{\mathcal{P}}_{(T-1)}\right]_{ki} \alpha^{T-1}(k)}{\sum_{i'=1}^{I} \sum_{k=1}^{I} \hat{\mathcal{E}}(i', \mathbf{y}^{\tau_T}) \left[\hat{\mathcal{P}}_{(T-1)}\right]_{ki'} \alpha^{T-1}(k)}
\end{aligned} \tag{3.35}
$$

Now, let's focus on calculating $\mathbb{E}[\Omega^{\mathbf{z}}|\mathbf{z}^{\tau_T} = i, \tilde{\mathbf{u}}, w, \hat{\Theta}]$. The generator matrix associated

with applying intervention $\tilde{\mathbf{u}}$ on a given patient with covariates w can be expressed as $[\hat{\mathcal{Q}}(\tilde{\mathbf{u}})]_{i\tilde{i}} = \hat{\delta}_{i\tilde{i}\,\tilde{\mathbf{u}}} \exp(\hat{\rho}'_{i\tilde{i}\,\tilde{\mathbf{u}}} w)$, where i and $\tilde{i}$ are the initial and final health states, respectively. This generator matrix can be used for estimating the transition probability function of the CT-HMM as $\mathcal{P}_{\tilde{\mathbf{u}}}(\Delta\tau) = \exp(\Delta\tau \hat{\mathcal{Q}}(\tilde{\mathbf{u}})) = \sum_{m=0}^{\infty} (\Delta\tau)^m \frac{\hat{\mathcal{Q}}(\tilde{\mathbf{u}})^m}{m!}$. Starting from state $\mathbf{z}^{\tau_T} = i$, the probability distribution of the underlying health state at the end of planning period will then be $\mathbb{P}(\tilde{\mathbf{z}} = \tilde{i} \,|\, \mathbf{z}^{\tau_T} = i, \tilde{\mathbf{u}}, w, \hat{\Theta}) = \left[\mathcal{P}_{\tilde{\mathbf{u}}}(\Delta\tau)\right]_{i\tilde{i}}$. For any given value of the underlying health state $\tilde{\mathbf{z}} = \tilde{i}$, we approximate $\Omega^{\mathbf{z}}$ with a known constant $\Omega^{\mathbf{z}}_{\tilde{i}} = \mathbb{E}[\Omega^{\mathbf{z}} | \tilde{\mathbf{z}} = \tilde{i}]$. Hence, the above expectation can be expressed as:

$$\mathbb{E}[\Omega^{\mathbf{z}} | \mathbf{z}^{\tau_T} = i, \tilde{\mathbf{u}}, w, \hat{\Theta}] = \sum_{\tilde{i}=1}^{I} \Omega^{\mathbf{z}}_{\tilde{i}} \, \mathbb{P}(\tilde{\mathbf{z}} = \tilde{i} \,|\, \mathbf{z}^{\tau_T} = i, \tilde{\mathbf{u}}, w, \hat{\Theta}) = \sum_{\tilde{i}=1}^{I} \Omega^{\mathbf{z}}_{\tilde{i}} \left[\mathcal{P}_{\tilde{\mathbf{u}}}(\Delta\tau)\right]_{i\tilde{i}} \tag{3.36}$$

where w implicitly affects $\mathcal{P}_{\tilde{\mathbf{u}}}(\Delta\tau)$. Replacing Equation 3.36 back in Equation 3.35 will give us a simplified representation of the expected disutility given y and w:

$$\mathbb{E}[\Omega^{\mathbf{z}} | \tilde{\mathbf{u}}, \mathbf{y}^{\tau_{1:T}}, \mathbf{u}^{\tau_{1:T-1}}, w, \hat{\Theta}] = \frac{\sum_{i=1}^{I} \sum_{k=1}^{I} \sum_{\tilde{i}=1}^{I} \Omega^{\mathbf{z}}_{\tilde{i}} \left[\mathcal{P}_{\tilde{\mathbf{u}}}(\Delta\tau)\right]_{i\tilde{i}} \hat{\mathcal{E}}(i, \mathbf{y}^{\tau_T}) \left[\hat{\mathcal{P}}_{(T-1)}\right]_{ki} \alpha^{T-1}(k)}{\sum_{i'=1}^{I} \sum_{k=1}^{I} \hat{\mathcal{E}}(i', \mathbf{y}^{\tau_T}) \left[\hat{\mathcal{P}}_{(T-1)}\right]_{ki'} \alpha^{T-1}(k)}$$

As discussed earlier, choosing the intervention option $\tilde{\mathbf{u}} = l$ contributes a component $\Omega^{\mathbf{u}}_l$ to the disutility due to its additional risk of experiencing side effects. The optimal decision, given the limited knowledge that we have about $\mathbf{z}^{\tau_T}$ based on the observed history, would then be:

$$\tilde{\mathbf{u}}^{hist} = \underset{0 \le l \le L-1}{argmin} \;\; \Omega^{\mathbf{u}}_l + \sum_{i=1}^{I} \sum_{k=1}^{I} \sum_{\tilde{i}=1}^{I} \Omega^{\mathbf{z}}_{\tilde{i}} \left[\mathcal{P}_l(\Delta\tau)\right]_{i\tilde{i}} \hat{\mathcal{E}}(i, \mathbf{y}^{\tau_T}) \left[\hat{\mathcal{P}}_{(T-1)}\right]_{ki} \alpha^{T-1}(k) \tag{3.37}$$

which can be calculated efficiently.

To quantify the impact of using the entire history of patient for decision making, we compare it with two alternative cases. If we had perfect knowledge about the underlying health state at the current period, i.e. $\mathbf{z}^{\tau_T} = i$ for some known i, then we could've chosen

the optimal intervention as $\tilde{\mathbf{u}}^* = \underset{0\leq l\leq L-1}{argmin}\ \ \Omega^{\mathbf{u}}_l + \sum_{\tilde{i}=1}^{I} \Omega^{\mathbf{z}}_{\tilde{i}} \left[\mathcal{P}_l(\Delta\tau)\right]_{i\tilde{i}}$. On the other hand, we consider a case in which the only piece of information that is used for decision making is the physician's observation at the current period $\mathbf{y}^{\tau_T}$ (rather than the entire history). In this scenario, the posterior probability of $\mathbf{z}^{\tau_T}$ is $\mathbb{P}(\mathbf{z}^{\tau_T} = i \,|\, \mathbf{y}^{\tau_T}, w, \hat{\Theta}) = \frac{\hat{\mathcal{E}}(i,\mathbf{y}^{\tau_T})\,\hat{\pi}_i}{\sum_{i'=1}^{I} \hat{\mathcal{E}}(i',\mathbf{y}^{\tau_T})\,\hat{\pi}_{i'}}$. Consequently, the optimal decision would be:

$$\tilde{\mathbf{u}}^{last} = \underset{0\leq l\leq L-1}{argmin}\ \ \Omega^{\mathbf{u}}_l + \sum_{i=1}^{I}\sum_{\tilde{i}=1}^{I} \Omega^{\mathbf{z}}_{\tilde{i}} \left[\mathcal{P}_l(\Delta\tau)\right]_{i\tilde{i}} \hat{\mathcal{E}}(i,\mathbf{y}^{\tau_T})\,\hat{\pi}_i \tag{3.38}$$

3.7.2 Numerical Experimentation on Synthetic Data

To quantify the impact of using the entire patient history for planning, we use Monte Carlo simulation. Specifically, we generate a dataset of training patient records that is used for estimating the model parameters. Afterward, we generate a test dataset that contains historical records for a set of patients. We then compare the single-period sub-optimality of the aforementioned two approaches versus the optimal planning approach in which we have full information about the underlying health condition at the planning time. All the codes are publicly available on *GitHub*[3]. We present the details regarding the choice of parameters in the appendix.

First, we compare the aforementioned two approaches based on their estimated distribution of the current underlying health state. In particular, we're interested to investigate the added value of having access to the entire patient history for estimating the underlying health state at the current time period. Therefore, for each patient in the test set we estimate the distribution of $\mathbf{z}^{\tau_T}$ once using the entire history and once using only the last physician observation data. For each one of these two cases, we take the state with the highest probability as the predicted state and compare it with the ground truth. By averaging over the entire test set, we obtain

[3]Refer to the Planning directory in *https://github.com/saeedghodsi93/Disease_Progression_Modeling_HMM*

the following normalized confusion matrix for the predicted $\mathbf{z}^{\tau_T}$:

$$\text{"entire history"}: \begin{bmatrix} 0.11 & 0.13 & 0.03 \\ 0.06 & 0.25 & 0.08 \\ 0.01 & 0.09 & 0.24 \end{bmatrix}, \qquad \text{"last observation"}: \begin{bmatrix} 0.18 & 0.00 & 0.09 \\ 0.14 & 0.00 & 0.25 \\ 0.01 & 0.03 & 0.30 \end{bmatrix}$$

where the rows correspond to the true value of $\mathbf{z}^{\tau_T}$ (up to down) and the columns correspond to the predicted value (left to right). Although it seems that using the entire history decreases the prediction accuracy for states 1 and 3, we observe that using only the last physician observation data makes prediction of state 2 extremely inaccurate. The overall prediction accuracy for the two cases are 60% and 48%, respectively.

As our measure of the optimal (i.e. full-information) disutility, for each patient in the test set we find $\tilde{\mathbf{u}}^*$ and use the optimal objective value $\Omega^* = \Omega^{\mathbf{u}}_{\tilde{\mathbf{u}}^*} + \sum_{\tilde{i}=1}^{I} \Omega^{\mathbf{z}}_{\tilde{i}} \left[\mathcal{P}_{\tilde{\mathbf{u}}^*}(\Delta\tau)\right]_{i\tilde{i}}$ (where $i = \mathbf{z}^{\tau_T}$ is known) as our target disutility. On the other hand, we find $\tilde{\mathbf{u}}^{hist}$ and $\tilde{\mathbf{u}}^{last}$ according to equations 3.37 and 3.38, respectively, and calculate the corresponding actual expected disutilities as $\Omega^{hist} = \Omega^{\mathbf{u}}_{\tilde{\mathbf{u}}^{hist}} + \sum_{\tilde{i}=1}^{I} \Omega^{\mathbf{z}}_{\tilde{i}} \left[\mathcal{P}_{\tilde{\mathbf{u}}^{hist}}(\Delta\tau)\right]_{i\tilde{i}}$ and $\Omega^{last} = \Omega^{\mathbf{u}}_{\tilde{\mathbf{u}}^{last}} + \sum_{\tilde{i}=1}^{I} \Omega^{\mathbf{z}}_{\tilde{i}} \left[\mathcal{P}_{\tilde{\mathbf{u}}^{last}}(\Delta\tau)\right]_{i\tilde{i}}$ (where $i = \mathbf{z}^{\tau_T}$ is the true state). To evaluate the effect of having access to more data on the decision-making outcome, we first compare the decisions $\tilde{\mathbf{u}}^{hist}$ and $\tilde{\mathbf{u}}^{last}$ with $\tilde{\mathbf{u}}^*$:

$$\text{"entire history"}: \begin{bmatrix} 0.15 & 0.08 & 0.01 \\ 0.09 & 0.20 & 0.11 \\ 0.02 & 0.05 & 0.29 \end{bmatrix}, \qquad \text{"last observation"}: \begin{bmatrix} 0.16 & 0.0 & 0.08 \\ 0.13 & 0.0 & 0.27 \\ 0.01 & 0.02 & 0.33 \end{bmatrix}$$

Here rows correspond to $\tilde{\mathbf{u}}^*$ (up to down) and columns correspond to $\tilde{\mathbf{u}}^{hist}$ and $\tilde{\mathbf{u}}^{last}$ (left to right). As the confusion matrices indicate, using the entire history leads to choosing intervention 3 with a higher probability whenever it is the optimal intervention. The overall accuracy of making the optimal decision is 64% and 49%, respectively, for these two cases. Notice that we've intentionally designed the true model in a way that choosing the optimal intervention requires making an accurate estimation for the underlying health

state. Specifically, the the optimal choice is $l = 0$, $l = 1$, and $l = 2$ when $\mathbf{z}^{\tau_T} = 1$, $\mathbf{z}^{\tau_T} = 2$, and $\mathbf{z}^{\tau_T} = 3$, respectively. Finally, we calculate the disutility sub-optimality values as $(\Omega^{hist} - \Omega^*)/\Omega^* = 7.3\%$ and $(\Omega^{last} - \Omega^*)/\Omega^* = 10.6\%$. In other words, using the entire history as a proxy for estimating the underlying health state increases the disutility by 7.3%, while using only the last physician observation increases the disutility to 10.6%.

3.7.3 Multi-Period Planning

In this section, we present a multi-period formulation for the planning problem. In particular, we assume the parameter set $\hat{\Theta} = (\hat{\pi}, \hat{\rho}, \hat{\delta}, \hat{\mu}, \hat{\eta})$ has been estimated using a training dataset. Moreover, we assume access to the history of physician observations $\mathbf{y}^{\tau_{1:T}}$ as well as the prescribed interventions $\mathbf{u}^{\tau_{1:T-1}}$. We then formulate an infinite-horizon reinforcement learning planning problem by estimating the environment parameters based on $\hat{\Theta}$ to determine $\tilde{\mathbf{u}}^s$ $(s = 1, 2, \cdots)$ in a way that the total expected discountedw disutility is minimized. Similar to the population-level planning section, here we follow a model predictive control approach as well, i.e. only the first action is executed and the optimization problem is solved again at the beginning of the next period when new data becomes available. Since the underlying health condition is assumed to be unobserved, we use a discrete-time partially-observed MDP model for planning. We discuss how traditional POMDP solver algorithms can be used for finding the optimal course of interventions.

The underlying health state, physician observation, and intervention are called "state", "observation", and "action", respectively, in the terminology of the POMDP literature. The "reward" $\Omega_s(\tilde{\mathbf{u}}^s, \tilde{\mathbf{z}}^s)$ associated with each planning period can be broken down to the sum of $\Omega^{\mathbf{u}}_{\tilde{\mathbf{u}}^s}$ (i.e. the intervention component) and $\Omega^{\mathbf{z}}_{\tilde{\mathbf{z}}^s}$ (i.e. the health state component) respectively. The total discounted reward would then be $\Omega = \sum_{s=1}^{\infty} \varrho^{s-1}\, \Omega_s$, where ϱ is the discount factor. The initial state probability (i.e. $\mathbb{P}(\tilde{\mathbf{z}}^0) = \mathbb{P}(\mathbf{z}^{\tau_T})$) can be determined according to Equation 3.34 based on the patients historical data. The transition matrix of the underlying health state, given intervention $\tilde{\mathbf{u}}$ at any arbitrary time period, will also be $\mathcal{P}_{\tilde{\mathbf{u}}}(\Delta\tau) = \exp(\Delta\tau \hat{\mathcal{Q}}(\tilde{\mathbf{u}})) =$

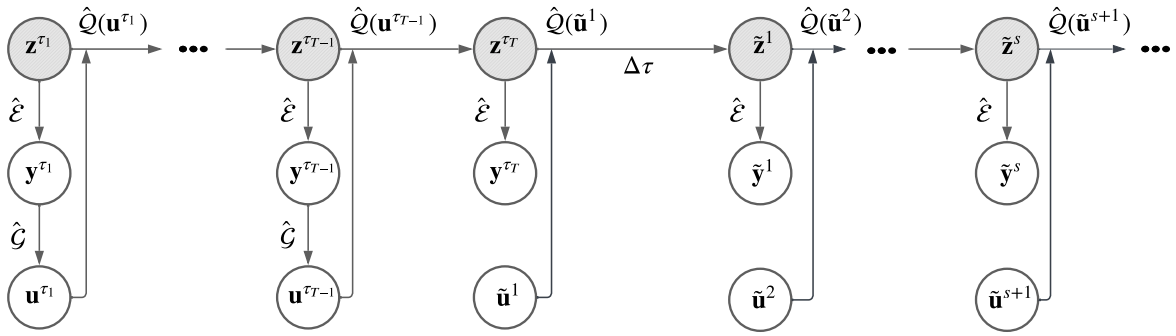

Figure 3.8: Multi-period planning, based on the patient's history.

$\sum_{m=0}^{\infty}(\Delta\tau)^m \frac{\hat{\mathcal{Q}}(\tilde{\mathbf{u}})^m}{m!}$, where $\hat{\mathcal{Q}}(\tilde{\mathbf{u}})$ is the estimated generator matrix. We denote the transition kernel by $\mathcal{T}$ for future reference. Finally, the observation distribution is given by $\hat{\mathcal{E}}$.

For POMDPs, the belief distribution over the latent state, i.e. $b^s(i) \coloneqq \mathbb{P}(\tilde{\mathbf{z}}^s = i|\tilde{\mathbf{y}}^{1:s}, \tilde{\mathbf{u}}^{1:s})$ $(1 \leq i \leq I)$, is known to be a sufficient statistic [Sag18]. Thus, a deterministic intervention policy can be formulated as a mapping from the space of belief distribution over the current latent health state to the space of interventions, i.e. $\Pi : \Delta^I \to \{0, \cdots, L-1\}$, where Δ^I is the I-dimensional probability simplex. In the following, we consider an infinite-horizon Point-Based Value Iteration (PBVI) algorithm [PGT03] for solving our POMDP. We refer the readers to [SPK13] for a comprehensive review of different variations of the PBVI algorithm and to [WS19] for a finite-horizon version of the algorithm. Define the belief-action value function associated with policy Π as:

$$Q(b, l; \Pi) = \mathbb{E}\big[\sum_{s=1}^{\infty} \varrho^{s-1}\, \Omega(\tilde{\mathbf{u}}^s, \tilde{\mathbf{z}}^s) \,|\, b^0 = b, \tilde{\mathbf{u}}^1 = l\big]$$

where the interventions are determined according to $\tilde{\mathbf{u}}^{s+1} = \Pi(b^s)$ for $s \geq 1$. The optimal policy would then be $\Pi^* = \arg\max_{\Pi} Q(b, l; \Pi)$. Denote by Q^* the optimal belief-action value function, i.e. $Q^*(b, l) = Q(b, l; \Pi^*)$. Furthermore, define the optimal belief value function as

$V^*(b) = \max_l Q^*(b,l)$. The Bellman update equation can be expressed as:

$$Q^*(b,l) = \langle \tilde{\Omega}_l, b \rangle + \varrho \sum_{j=0}^{J-1} \mathbb{P}(\tilde{\mathbf{y}}' = j \,|\, b, l)\, V^*(\tilde{b})$$

where $\tilde{\Omega}_l = \Omega_s(l,\cdot)$ and $\tilde{b}(k) \coloneqq \mathbb{P}(k \,|\, b, i, j)$ $(1 \leq k \leq I)$.

It is known that $V^*(b)$ is a piecewise-linear convex function of b [Son78]. Specifically, the value function can be represented using a set of vectors Γ as $V^*(b) = \max_{\alpha \in \Gamma} \langle \alpha, b \rangle$. Exact Value Iteration algorithms update the set Γ at each iteration by running the complete Bellman backup procedure. It turns out that the set Γ grows exponentially with each iteration in exact Value Iteration algorithms. Hence, PBVI algorithms try to make the procedure computationally tractable by evaluating the Bellman backup on only a finite subset $\mathcal{S} \subset \Delta^I$ of belief values. In particular, let's fix $b \in \mathcal{S}$ and define a transformed version of Γ for each l and j as $\tilde{\Gamma}_{lj} \coloneqq \{\tilde{\alpha} \,|\, \tilde{\alpha}(i) = \sum_{k=1}^{I} \alpha(k)\, \mathcal{T}_{il}(k)\, \mathcal{E}(k,j);\, \forall\, 1 \leq i \leq I,\, \alpha \in \Gamma\}$, where $\mathcal{T}$ is the transition kernel and $\mathcal{E}$ is the emission distribution. After some algebra, the Bellman update equation can be written as:

$$Q^*(b,l) = \langle \tilde{\Omega}_l, b \rangle + \varrho \sum_{j=0}^{J-1} \max_{\tilde{\alpha} \in \tilde{\Gamma}_{lj}} \langle \tilde{\alpha}, b \rangle$$

Let $\alpha_l^*(b;\Gamma) \coloneqq \tilde{\Omega}_l + \varrho \sum_{j=0}^{J-1} \arg\max_{\tilde{\alpha} \in \tilde{\Gamma}_{lj}} \langle \tilde{\alpha}, b \rangle$ and observe that $Q^*(b,l) = \langle \alpha_l^*(b;\Gamma), b \rangle$. Therefore, the Bellman backup at belief b can be constructed as

$$\alpha^*(b;\Gamma) \coloneqq \arg\max_{\{\alpha_l^*(b;\Gamma) \,|\, 0 \leq l \leq L-1\}} \langle \alpha_a^*(b;\Gamma), b \rangle$$

To be more specific, had we ran the complete Bellman backup on Γ, then $\alpha^*(b;\Gamma)$ would have been the value-maximizing vector in the resulting set of vectors at belief b. PBVI algorithms start by choosing an initial set of belief points $\mathcal{S}^{(0)}$ as well as an initial value function set of vectors $\Gamma^{(0)}$ and update these sets iteratively. At iteration n, point-based Bellman backup is

performed for updating the value function set of vectors as $\Gamma^{(n+1)} = \{\alpha^*(b; \Gamma^{(n)}) \,|\, b \in \mathcal{S}^{(n)}\}$. Many different approaches have been proposed for constructing the set $\mathcal{S}^{(n+1)}$ based on $\mathcal{S}^{(n)}$. For example, one strategy would be to choose beliefs evenly spaced on a grid. Alternatively, one can add belief points that are farthest to the current set $\mathcal{S}^{(n)}$. Following any one of these strategies, we can update the sets and find the optimal course of interventions.

3.8 Conclusion

In this chapter, we studied the problem of data-driven healthcare intervention planning for both behavioral and medical healthcare sectors. Specifically, we provided a discussion on the inefficiency issues in the US healthcare system and reviewed the existing literature from the economic and public health perspectives. Our research was originally motivated by observing the trade-off between providing care as needed versus early intervention, which is common especially within the behavioral healthcare sector. Therefore, we formulated a population-level intervention planning problem for providing guidelines regarding assignment of intervention packages to patient groups. The main challenge in data-driven planning is in appropriate estimation of the transition rates based on historical data. To do so, we developed a statistical disease progression model that describes the continuous-time process of disease using discrete-time irregularly-spaced observations. Our main contribution here is to present an efficient EM algorithm for estimating the parameters of the CT-HMM model. We generated synthetic data and provided a numerical analysis for the convergence of parameters to the true values.

In an extended version of the model, we assumed the underlying health state is occasionally observable. The motivation for this model is the fact that for certain medical domains accurate examination is possible, but often expensive and risky. For example, in caner maintenance therapy the physician can order biopsy to examine the true underlying health condition of the patient, but it is done only in certain circumstances. Our extended model also incorporates

the direct effect of interventions on physician observations. For the case of cancer maintenance therapy, some interventions such as medication play the role of symptom alleviation rather than improving the underlying health condition. Therefore, we added a direct link from interventions to physician observations in the model. We developed an EM algorithm for efficiently estimating the parameters and demonstrated the convergence of our algorithm using synthetically generated data. Finally, we analyzed the individual-level intervention planning problem in both the single-period and multi-period settings.

3.9 Appendix A: Baseline Model

3.9.1 Simulation Setup

In this section, we provide more details regarding the numerical experiments of the baseline model. We refer the readers to [Gho22] for more details. Specifically, we set the time parameters of the trajectories as $\mu_T = 10$, $\sigma_T = 3$, $\mu_\tau = 10$, $\sigma_\tau = 3$. Moreover, we set the age parameters as $\mu_{a,y} = 25$, $\sigma_{a,y} = 5$, $\mu_{a,o} = 50$, $\sigma_{a,o} = 10$, $p_a = 0.25$. The transition parameters of the underlying CTMC are:

$$\zeta^*_{\cdot,\cdot,0} = \begin{bmatrix} 0.00 & 0.60 & 0.40 \\ 0.20 & 0.00 & 0.80 \\ 0.10 & 0.90 & 0.00 \end{bmatrix}, \qquad \zeta^*_{\cdot,\cdot,1} = \begin{bmatrix} 0.00 & 0.80 & 0.20 \\ 0.60 & 0.00 & 0.40 \\ 0.30 & 0.70 & 0.00 \end{bmatrix}, \qquad \zeta^*_{\cdot,\cdot,2} = \begin{bmatrix} 0.00 & 0.90 & 0.10 \\ 0.80 & 0.00 & 0.20 \\ 0.50 & 0.50 & 0.00 \end{bmatrix}$$

where the last dimension corresponds to the intervention. The mean time spent in each state are also set according to:

$$\iota^*_{\cdot,0} = \begin{bmatrix} 4 & 6 & 20 \end{bmatrix}, \qquad \iota^*_{\cdot,1} = \begin{bmatrix} 9 & 12 & 10 \end{bmatrix}, \qquad \iota^*_{\cdot,2} = \begin{bmatrix} 11 & 10 & 7 \end{bmatrix}$$

Hence, the true generator matrix parameters are:

$$\delta^*_{\cdot,\cdot,0} = \begin{bmatrix} - & 0.15 & 0.10 \\ 0.03 & - & 0.13 \\ 0.00 & 0.04 & - \end{bmatrix}, \quad \delta^*_{\cdot,\cdot,1} = \begin{bmatrix} - & 0.09 & 0.02 \\ 0.05 & - & 0.03 \\ 0.03 & 0.07 & - \end{bmatrix}, \quad \delta^*_{\cdot,\cdot,2} = \begin{bmatrix} - & 0.08 & 0.01 \\ 0.08 & - & 0.02 \\ 0.07 & 0.07 & - \end{bmatrix}$$

Moreover, the age coefficients are set to:

$$\rho^*_{\cdot,\cdot,0} = \rho^*_{\cdot,\cdot,1} = \rho^*_{\cdot,\cdot,2} = \begin{bmatrix} - & 0.020 & 0.020 \\ 0.015 & - & 0.015 \\ 0.010 & 0.010 & - \end{bmatrix}$$

Assuming the age of an individual is 50 years, we'll get the following parameters:

$$\mathcal{Q}^*_{\cdot,\cdot,0} = \begin{bmatrix} -0.68 & 0.41 & 0.27 \\ 0.07 & -0.35 & 0.28 \\ 0.01 & 0.07 & -0.08 \end{bmatrix}, \qquad \mathcal{Q}^*_{\cdot,\cdot,1} = \begin{bmatrix} -0.30 & 0.24 & 0.06 \\ 0.11 & -0.18 & 0.07 \\ 0.05 & 0.12 & -0.16 \end{bmatrix}$$

$$\mathcal{Q}^*_{\cdot,\cdot,2} = \begin{bmatrix} -0.25 & 0.22 & 0.02 \\ 0.17 & -0.21 & 0.04 \\ 0.12 & 0.12 & -0.24 \end{bmatrix}$$

The true initial health state probability is set to $\pi^* = \left[expm(1000 \times \mathcal{Q}^*_{\cdot,\cdot,0})\right]_{0,\cdot} = (0.03, 0.19, 0.77)$. The reason behind this choice is that it makes the initial state probability equal to the stationary distribution of the no-intervention option. Finally, the emission and intervention probabilities are set to $\mu^* = (0.1, 0.5, 0.9)$ and

$$\eta^* = (0.04, 0.15, 0.26, 0.32, 0.43, 0.51, 0.62, 0.77, 0.81, 0.90)$$

For the starting value of the initial health state parameter, we chose $\tilde{\pi} = (0.33, 0.33, 0.33)$. The corresponding parameters of the underlying CTMC are also set to:

$$\tilde{\zeta}_{\cdot,\cdot,0} = \begin{bmatrix} 0.00 & 0.50 & 0.50 \\ 0.30 & 0.00 & 0.70 \\ 0.20 & 0.80 & 0.00 \end{bmatrix}, \qquad \tilde{\zeta}_{\cdot,\cdot,1} = \begin{bmatrix} 0.00 & 0.60 & 0.40 \\ 0.40 & 0.00 & 0.60 \\ 0.40 & 0.60 & 0.00 \end{bmatrix}, \qquad \tilde{\zeta}_{\cdot,\cdot,2} = \begin{bmatrix} 0.00 & 0.70 & 0.30 \\ 0.60 & 0.00 & 0.40 \\ 0.40 & 0.60 & 0.00 \end{bmatrix}$$

$$\tilde{\iota}_{\cdot,0} = \begin{bmatrix} 5 & 8 & 12 \end{bmatrix}, \qquad \tilde{\iota}_{\cdot,1} = \begin{bmatrix} 6 & 10 & 10 \end{bmatrix}, \qquad \tilde{\iota}_{\cdot,2} = \begin{bmatrix} 9 & 8 & 4 \end{bmatrix}$$

We'll get the following starting values for the generator matrix parameters:

$$\tilde{\delta}_{\cdot,\cdot,0} = \begin{bmatrix} - & 0.10 & 0.10 \\ 0.04 & - & 0.09 \\ 0.02 & 0.07 & - \end{bmatrix}, \quad \tilde{\delta}_{\cdot,\cdot,1} = \begin{bmatrix} - & 0.10 & 0.07 \\ 0.04 & - & 0.06 \\ 0.04 & 0.06 & - \end{bmatrix}, \quad \tilde{\delta}_{\cdot,\cdot,2} = \begin{bmatrix} - & 0.08 & 0.03 \\ 0.07 & - & 0.05 \\ 0.10 & 0.15 & - \end{bmatrix}$$

For the age parameters, we set:

$$\tilde{\rho}_{\cdot,\cdot,0} = \tilde{\rho}_{\cdot,\cdot,1} = \tilde{\rho}_{\cdot,\cdot,2} = \begin{bmatrix} - & 0.01 & 0.01 \\ 0.01 & - & 0.01 \\ 0.01 & 0.01 & - \end{bmatrix}$$

Moreover, the emission and intervention parameters are chosen as $\tilde{\mu} = (0.45, 0.65, 0.7)$ and $\tilde{\eta} = (0.01, 0.03, 0.06, 0.08, 0.45, 0.48, 0.51, 0.55, 0.91, 0.97)$.

3.9.2 An Overview of the Algorithm

Here we provide an end-to-end description of the estimation algorithm. For calculating the conditional expectations corresponding to each interval $(\tau_{ns}, \tau_{n\,s+1}]$, we first construct $\tilde{\mathcal{Q}}_{n\,(s)}$ based on 3.3, 3.4. Afterward, the integral $\xi_{ik} = \int_0^{\tau_{n\,s+1}-\tau_{ns}} \exp\left(x\tilde{\mathcal{Q}}_{n\,(s)}\right) \mathcal{B}_{ik} \exp\left((\tau_{n\,s+1} - \tau_{ns} - x)\tilde{\mathcal{Q}}_{n\,(s)}\right) dx$ is calculated for all (i, k). To do so, we evaluate $\exp\left((\tau_{n\,s+1} - \tau_{ns})\mathcal{A}_{ik}\right)$ for $\mathcal{A}_{ik} = \begin{pmatrix} \tilde{\mathcal{Q}}_{n\,(s)} & \mathcal{B}_{ik} \\ 0 & \tilde{\mathcal{Q}}_{n\,(s)} \end{pmatrix}$ and the upper right part of the result gives us the matrix that we're looking for. For $\tilde{i}$ and $\tilde{k}$ being the values of the end states, we can find the (i, k)-th expectation of 3.1, 3.2 in the $(\tilde{i}, \tilde{k})$-th element of the matrix.

Algorithm 3: The estimation algorithm

Input: $\tilde{\Theta} = (\tilde{\pi}, \tilde{\delta}, \tilde{\rho}, \tilde{\mu}, \tilde{\eta})$
Output: $\hat{\Theta} = (\hat{\pi}, \hat{\delta}, \hat{\rho}, \hat{\mu}, \hat{\eta})$
Data: $(\mathbf{y}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}})\big|_{1\le s\le T_n}$ for $1 \le n \le N$
while not converged **do**

E-step: for $1 \le n \le N$

- Calculate $\alpha_n^s(i)$ based on 3.9, 3.11 and $\beta_n^s(i)$ based on 3.10, 3.12
- Calculate $\gamma_n^s(i) = \frac{\alpha_n^s(i)\beta_n^s(i)}{\sum_{i'=1}^{I}\alpha_n^s(i')\beta_n^s(i')}$ and $\nu_n^s(i,k)$ using 3.13
- Calculate $\mathbb{E}[\chi_{n\,ik}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta})]$ and $\mathbb{E}[\psi_{n\,i}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}]$

M-step:

- Set $\tilde{\pi}_i = \frac{\sum_{n=1}^{N}\gamma_n^1(i)}{\sum_{i'=1}^{I}\sum_{n=1}^{N}\gamma_n^1(i')}$, $\tilde{\mu}_i = \frac{\sum_{n=1}^{N}\sum_{s=1}^{T_n}\mathbf{y}_n^{\tau_{ns}}\gamma_n^s(i)}{(J-1)\sum_{n=1}^{N}\sum_{s=1}^{T_n}\gamma_n^s(i)}$, $\tilde{\eta}_j = \frac{\sum_{n=1}^{N}\sum_{s=1}^{T_n}\mathbf{u}_n^{\tau_{ns}}\mathbb{I}\{\mathbf{y}_n^{\tau_{ns}}=j\}}{(L-1)\sum_{n=1}^{N}\sum_{s=1}^{T_n}\mathbb{I}\{\mathbf{y}_n^{\tau_{ns}}=j\}}$
- Update $\tilde{\delta}_{ikl}$ based on 3.6 and $\tilde{\rho}_{ikl}$ based on 3.7, 3.8

return $\hat{\pi} = \tilde{\pi}$, $\hat{\delta} = \tilde{\delta}$, $\hat{\rho} = \tilde{\rho}$, $\hat{\mu} = \tilde{\mu}$, $\hat{\eta} = \tilde{\eta}$

3.10 Appendix B: Intervention Planning

3.10.1 Implementation Details

The details of implementation for the planning section are very similar to those of the baseline model. In the following, we'll present the set of parameters that we used for generating the training and test datasets. We experiment with different values for the training set size N_{train}, but fix the test size at $N_{test} = 10,000$.

To set the generator matrix, we construct the transition probability matrix of the embedded DTMC in the following way:

$$\zeta_{\cdot,\cdot,0}^* = \begin{bmatrix} 0.00 & 0.90 & 0.10 \\ 0.20 & 0.00 & 0.80 \\ 0.10 & 0.90 & 0.00 \end{bmatrix}, \quad \zeta_{\cdot,\cdot,1}^* = \begin{bmatrix} 0.00 & 0.70 & 0.30 \\ 0.65 & 0.00 & 0.35 \\ 0.15 & 0.85 & 0.00 \end{bmatrix}, \quad \zeta_{\cdot,\cdot,2}^* = \begin{bmatrix} 0.00 & 0.55 & 0.45 \\ 0.50 & 0.00 & 0.50 \\ 0.60 & 0.40 & 0.00 \end{bmatrix}$$

The mean sojourn times associated with each of the states are defined as:

$$\iota^*_{\cdot,0} = \begin{bmatrix} 18 & 5 & 25 \end{bmatrix}, \qquad \iota^*_{\cdot,1} = \begin{bmatrix} 9 & 11 & 10 \end{bmatrix}, \qquad \iota^*_{\cdot,2} = \begin{bmatrix} 4 & 8 & 6 \end{bmatrix}$$

Consequently, the true baseline generator matrix elements can then be constructed as:

$$\delta^*_{\cdot,\cdot,0} = \begin{bmatrix} - & 0.05 & 0.01 \\ 0.04 & - & 0.16 \\ 0.00 & 0.04 & - \end{bmatrix}, \quad \delta^*_{\cdot,\cdot,1} = \begin{bmatrix} - & 0.08 & 0.03 \\ 0.06 & - & 0.03 \\ 0.01 & 0.08 & - \end{bmatrix}, \quad \delta^*_{\cdot,\cdot,2} = \begin{bmatrix} - & 0.14 & 0.11 \\ 0.06 & - & 0.06 \\ 0.10 & 0.07 & - \end{bmatrix}$$

For our 50-years old example patient, the true generator matrix will be:

$$\mathcal{Q}^*_{\cdot,\cdot,0} = \begin{bmatrix} -0.15 & 0.14 & 0.02 \\ 0.08 & -0.42 & 0.34 \\ 0.01 & 0.06 & -0.07 \end{bmatrix}, \qquad \mathcal{Q}^*_{\cdot,\cdot,1} = \begin{bmatrix} -0.30 & 0.21 & 0.09 \\ 0.13 & -0.19 & 0.07 \\ 0.02 & 0.14 & -0.16 \end{bmatrix}$$

$$\mathcal{Q}^*_{\cdot,\cdot,2} = \begin{bmatrix} -0.68 & 0.37 & 0.31 \\ 0.13 & -0.26 & 0.13 \\ 0.16 & 0.11 & -0.27 \end{bmatrix}$$

and the corresponding probability distribution of the initial state is $\pi^* = \left[expm(1000 \times \mathcal{Q}^*_{\cdot,\cdot,0})\right]_{0,\cdot} = (0.12, 0.14, 0.72)$. To make it difficult for the algorithm to extract information about the underlying health state from the physician observation, we choose $\mu^* = (0.45, 0.55, 0.75)$. As a consequence, states 1 and 2 would lool like relatively similar when we only consider the physician observation data. The other parameters are mostly the same as the baseline simulation. We choose the disutility coefficients associated with different underlying health states as $\Omega_{\cdot} = (2, 7, 13)$.

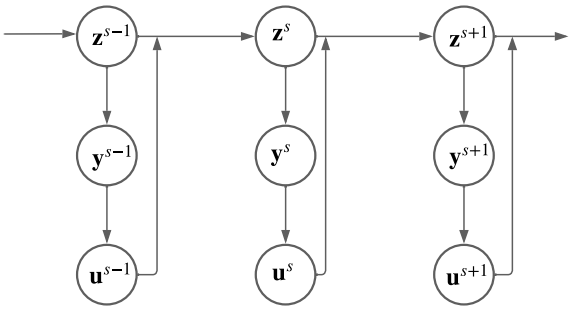

Figure 3.9: DT-HMM model with interventions at visit times.

3.11 Appendix C: Discrete-Time Model

3.11.1 MCEM Parameter Learning for a DT-HMM Disease Progression Model

As we discussed in Section 3.3.2, for simpler models such as DT-HMMs an alternative for calculating the aforementioned expectations and posterior probabilities would be to approximate them via the Monte Carlo approach. In particular, consider a DT-HMM disease progression model with interventions at visit times as in Figure 3.9. In this section, we discuss an MCEM parameter learning algorithm that approximates the posterior probabilities by generating samples from the posterior distribution of the latent variables of the graphical model based on the Gibbs sampling approach.

Similar to the baseline case, suppose we have N patients in the healthcare system and let T_n indicate the number of observations for patient n. In any time period $s \in \{1, \cdots, T_n\}$, denote by $\mathbf{z}_n^s$, $\mathbf{y}_n^s$, and $\mathbf{u}_n^s$ the underlying health state, physician observation, and intervention variables and suppose they are all discrete-valued as in the continuous-time case. The transition probability matrix of the underlying DTMC that governs the dynamics of process $\mathbf{z}$ between periods $(s, s+1]$ is denoted by $\mathcal{Q} \in \mathbb{R}^{I \times I \times L}$. Specifically, for the given intervention $\mathbf{u}_n^s = l$, the $I \times I$ matrix $[\mathcal{Q}]_{\cdot,\cdot,l}$ determines the transition rates of the underlying process between periods $(s, s+1]$. Needless to explains, we have $I(I-1)L$ free parameters for our transition rates matrix. Again, the emission distribution $\mathcal{E}(i,j) = \mathbb{P}(\mathbf{y}_n^s = j \,|\, \mathbf{z}_n^s = i)$ and

intervention distribution $\mathcal{G}(j,l) = \mathbb{P}(\mathbf{u}_n^s = l \,|\, \mathbf{y}_n^s = j)$ $(1 \leq s \leq T_n)$ are assumed to be binomial with parameters μ_i $(1 \leq i \leq I)$ and η_j $(0 \leq j \leq J-1)$, respectively. Finally, we denote by $\pi \in \mathbb{R}^I$ the vector of initial probabilities of $\mathbf{z}$.

Let $\Theta \triangleq (\pi, \mathcal{Q}, \mu, \eta)$. The complete-data log-likelihood for patient n can be expressed as:

$$\log\, \mathcal{L}_n(\Theta) = \log \mathbb{P}(\mathbf{z}_n^1) + \sum_{s=1}^{T_n-1} \big[\log \mathbb{P}(\mathbf{z}_n^{s+1}|\mathbf{z}_n^s, \mathbf{u}_n^s, \mathcal{Q})\big] + \sum_{s=1}^{T_n} \big[\log \mathbb{P}(\mathbf{u}_n^s|\mathbf{y}_n^s, \eta) + \log \mathbb{P}(\mathbf{y}_n^s|\mathbf{z}_n^s, \mu)\big]$$

For a given realization of $\mathbf{z}$, the function $\log\, \mathcal{L}_n(\Theta)$ can be simplified, in the following way:

$$\log \mathbb{P}(\mathbf{z}_n^{s+1}|\mathbf{z}_n^s, \mathbf{u}_n^s, \mathcal{Q}) = \sum_{i=1}^{I}\sum_{k=1}^{I}\sum_{l=0}^{L-1} \log[\mathcal{Q}]_{i,k,l}\, \mathbb{I}\{\mathbf{z}_n^s = i, \mathbf{z}_n^{s+1} = k\}\, \mathbb{I}\{\mathbf{u}_n^s = l\}$$

The complete-data log-likelihood for all the patients will then be:

$$\begin{aligned}
\log \mathcal{L}(\Theta) &= \sum_{n=1}^{N} \log\, \mathcal{L}_n(\Theta) = \sum_{n=1}^{N}\sum_{i=1}^{I} \log(\pi_i)\, \mathbb{I}\{\mathbf{z}_n^1 = i\} \\
&+ \sum_{n=1}^{N}\sum_{s=1}^{T_n-1}\sum_{i=1}^{I}\sum_{k=1}^{I}\sum_{l=0}^{L-1} \log[\mathcal{Q}]_{i,k,l}\, \mathbb{I}\{\mathbf{z}_n^s = i, \mathbf{z}_n^{s+1} = k\}\, \mathbb{I}\{\mathbf{u}_n^s = l\} \\
&+ \sum_{n=1}^{N}\sum_{s=1}^{T_n}\sum_{i=1}^{I} \left[\log \binom{J-1}{\mathbf{y}_n^s} + \mathbf{y}_n^s \log(\mu_i) + (J-1-\mathbf{y}_n^s)\log(1-\mu_i)\right] \mathbb{I}\{\mathbf{z}_n^s = i\} \\
&+ \sum_{n=1}^{N}\sum_{s=1}^{T_n}\sum_{j=0}^{J-1} \left[\log \binom{L-1}{\mathbf{u}_n^s} + \mathbf{u}_n^s \log(\eta_j) + (L-1-\mathbf{u}_n^s)\log(1-\eta_j)\right] \mathbb{I}\{\mathbf{y}_n^s = j\}
\end{aligned}$$

Define the initial parameters $\tilde{\Theta} \triangleq (\tilde{\pi}, \tilde{Q}, \tilde{\mu}, \tilde{\eta})$. The expected complete-data log-likelihood is:

$$\begin{aligned}
&\mathbb{E}[\log\, \mathcal{L}(\Theta)|\mathbf{u}_n^{1:T_n}, \mathbf{y}_n^{1:T_n}, \tilde{\Theta}] = \sum_{n=1}^{N}\sum_{i=1}^{I} \log(\pi_i)\, \mathbb{P}(\mathbf{z}_n^1 = i \,|\, \mathbf{u}_n^{1:T_n}, \mathbf{y}_n^{1:T_n}, \tilde{\Theta}) \\
&+ \sum_{n=1}^{N}\sum_{s=1}^{T_n-1}\sum_{i=1}^{I}\sum_{k=1}^{I}\sum_{l=0}^{L-1} \log[\mathcal{Q}]_{i,k,l}\, \mathbb{P}(\mathbf{z}_n^s = i, \mathbf{z}_n^{s+1} = k \,|\, \mathbf{u}_n^{1:T_n}, \mathbf{y}_n^{1:T_n}, \tilde{\Theta})\, \mathbb{I}\{\mathbf{u}_n^s = l\} \\
&+ \sum_{n=1}^{N}\sum_{s=1}^{T_n}\sum_{i=1}^{I} \left[\log \binom{J-1}{\mathbf{y}_n^s} + \mathbf{y}_n^s \log(\mu_i) + (J-1-\mathbf{y}_n^s)\log(1-\mu_i)\right] \mathbb{P}(\mathbf{z}_n^s = i \,|\, \mathbf{u}_n^{1:T_n}, \mathbf{y}_n^{1:T_n}, \tilde{\Theta})
\end{aligned}$$

$$+\sum_{n=1}^{N}\sum_{s=1}^{T_n}\sum_{j=0}^{J-1}\left[\log\binom{L-1}{\mathbf{u}_n^s}+\mathbf{u}_n^s\log(\eta_j)+(L-1-\mathbf{u}_n^s)\log(1-\eta_j)\right]\mathbb{I}\{\mathbf{y}_n^s=j\}$$

Define $\gamma_n^s(i)=\mathbb{P}(\mathbf{z}_n^s=i\,|\,\mathbf{u}_n^{1:T_n},\mathbf{y}_n^{1:T_n},\tilde{\Theta})$ (for $1\le s\le T_n$) and $\nu_n^s(i,k)=\mathbb{P}(\mathbf{z}_n^s=i,\mathbf{z}_n^{s+1}=k\,|\,\mathbf{u}_n^{1:T_n},\mathbf{y}_n^{1:T_n},\tilde{\Theta})$ (for $1\le s\le T_n-1$). In the M-step, we can update the parameters according to $\hat{\pi}_i=\frac{\sum_{n=1}^N\gamma_n^1(i)}{\sum_{i'=1}^I\sum_{n=1}^N\gamma_n^1(i')}$, $\hat{\mu}_i=\frac{\sum_{n=1}^N\sum_{s=1}^{T_n}\mathbf{y}_n^s\gamma_n^s(i)}{(J-1)\sum_{n=1}^N\sum_{s=1}^{T_n}\gamma_n^s(i)}$, $\hat{\eta}_j=\frac{\sum_{n=1}^N\sum_{s=1}^{T_n}\mathbf{u}_n^s\mathbb{I}\{\mathbf{y}_n^s=j\}}{(L-1)\sum_{n=1}^N\sum_{s=1}^{T_n}\mathbb{I}\{\mathbf{y}_n^s=j\}}$, and $[\hat{\mathcal{Q}}]_{i,k,l}=\frac{\sum_{n=1}^N\sum_{s=1}^{T_n-1}\nu_n^s(i,k)\,\mathbb{I}\{\mathbf{u}_n^s=l\}}{\sum_{\tilde{k}=1}^I\sum_{n=1}^N\sum_{s=1}^{T_n-1}\nu_n^s(i,\tilde{k})\,\mathbb{I}\{\mathbf{u}_n^s=l\}}$. Hence, the problem reduces to calculation or approximation of γ and ν. The direct approach for calculating these probabilities would be to use the forward-backward algorithm. On the other hand, we can approximate the probabilities by generating samples from the posterior of the latent variables. In the following, we present both of these approaches and compare their results.

To directly calculate these probabilities, we construct a time-inhomogeneous DT-HMM from our original DT-HMM, by defining the transition probability function of the DT-HMM to be equal to $[\tilde{\mathcal{P}}_{n\,(s)}]_{ik}=\mathbb{P}(\mathbf{z}_n^{s+1}=k|\mathbf{z}_n^s=i,\mathbf{u}_n^s,\tilde{\Theta})=\left[\tilde{Q}\right]_{i,k,\mathbf{u}_n^s}$ for the interval $(s,s+1]$. We start by calculating $\gamma_n^s(i)$. The forward and backward probabilities are defined as $\alpha_n^s(i)=\mathbb{P}(\mathbf{z}_n^s=i,\mathbf{u}_n^{1:s},\mathbf{y}_n^{1:s}|\tilde{\Theta})$ and $\beta_n^s(i)=\mathbb{P}(\mathbf{u}_n^{s+1:T_n},\mathbf{y}_n^{s+1:T_n}|\mathbf{z}_n^s=i,\mathbf{u}_n^s,\tilde{\Theta})$, respectively. The dynamic programming equations for calculating the forward and backward variables are:

$$\begin{aligned}
\alpha_n^s(i)&=\sum_{k=1}^{I}\tilde{\mathcal{G}}(\mathbf{y}_n^s,\mathbf{u}_n^s)\,\tilde{\mathcal{E}}(i,\mathbf{y}_n^s)\left[\tilde{\mathcal{P}}_{n\,(s-1)}\right]_{ki}\alpha_n^{s-1}(k)\\
\alpha_n^1(i)&=\tilde{\mathcal{G}}(\mathbf{y}_n^1,\mathbf{u}_n^1)\,\tilde{\mathcal{E}}(i,\mathbf{y}_n^1)\tilde{\pi}_i\\
\beta_n^s(i)&=\sum_{k=1}^{I}\tilde{\mathcal{G}}(\mathbf{y}_n^{s+1},\mathbf{u}_n^{s+1})\,\tilde{\mathcal{E}}(k,\mathbf{y}_n^{s+1})\left[\tilde{\mathcal{P}}_{n\,(s)}\right]_{ik}\beta_n^{s+1}(k)\\
\beta_n^{T_n-1}(i)&=\sum_{k=1}^{I}\tilde{\mathcal{G}}(\mathbf{y}_n^{T_n},\mathbf{u}_n^{T_n})\,\tilde{\mathcal{E}}(k,\mathbf{y}_n^{T_n})\left[\tilde{\mathcal{P}}_{n\,(T_n-1)}\right]_{ik}
\end{aligned}$$

The posterior probabilities γ and ν can then be calculated as $\gamma_n^s(i)=\frac{\alpha_n^s(i)\beta_n^s(i)}{\sum_{i'=1}^I\alpha_n^s(i')\beta_n^s(i')}$ and $\nu_n^s(i,k)=\frac{\tilde{\mathcal{G}}(\mathbf{y}_n^{s+1},\mathbf{u}_n^{s+1})\,\tilde{\mathcal{E}}(k,\mathbf{y}_n^{s+1})\left[\tilde{\mathcal{P}}_{n\,(s)}\right]_{ik}\alpha_n^s(i)\,\beta_n^{s+1}(k)}{\sum_{i'=1}^I\sum_{k'=1}^I\tilde{\mathcal{G}}(\mathbf{y}_n^{s+1},\mathbf{u}_n^{s+1})\,\tilde{\mathcal{E}}(k',\mathbf{y}_n^{s+1})\left[\tilde{\mathcal{P}}_{n\,(s)}\right]_{i'k'}\alpha_n^s(i')\,\beta_n^{s+1}(k')}$ using α and β.

Now, we present our alternative Monte Carlo approach for approximating the posterior probabilities. Specifically, we generate a set of samples from the joint distribution of the graphical model (i.e. $\mathbb{P}(\mathbf{u}_n^{1:T_n}, \mathbf{y}_n^{1:T_n}, \mathbf{z}_n^{1:T_n}|\tilde{\Theta})$) using to the Gibbs sampling strategy, and approximate the posterior probabilities $\gamma_n^s(\cdot)$, $\nu_n^s(\cdot,\cdot)$ by the corresponding sample averages. Since $\mathbf{u}_n^{1:T_n}$ and $\mathbf{y}_n^{1:T_n}$ are known, we essentially need to generate samples from $\mathbb{P}(\mathbf{z}_n^{1:T_n} \,|\, \mathbf{u}_n^{1:T_n}, \mathbf{y}_n^{1:T_n}, \tilde{\Theta})$. Hence, the sample generation procedure reduces to calculating the conditional distribution of each node $\mathbf{z}_n^s$ given the rest of the nodes as well as the current parameter estimate $\tilde{\Theta}$:

$$\begin{aligned}
&\mathbb{P}(\mathbf{z}_n^s = i \,|\, \mathbf{z}_n^{1:s-1}, \mathbf{z}_n^{s+1:T_n}, \mathbf{u}_n^{1:T_n}, \mathbf{y}_n^{1:T_n}, \tilde{\Theta}) = \mathbb{P}(\mathbf{z}_n^s = i \,|\, \mathbf{z}_n^{s-1}, \mathbf{z}_n^{s+1}, \mathbf{u}_n^{s-1}, \mathbf{u}_n^s, \mathbf{y}_n^s, \tilde{\Theta}) \\
&\quad \propto \mathbb{P}(\mathbf{z}_n^s = i, \mathbf{z}_n^{s+1}, \mathbf{u}_n^s, \mathbf{y}_n^s \,|\, \mathbf{z}_n^{s-1}, \mathbf{u}_n^{s-1}, \tilde{\Theta}) \\
&\quad = \mathbb{P}(\mathbf{z}_n^{s+1} \,|\, \mathbf{z}_n^s = i, \mathbf{u}_n^s, \tilde{\Theta})\, \mathbb{P}(\mathbf{u}_n^s \,|\, \mathbf{y}_n^s, \tilde{\Theta})\, \mathbb{P}(\mathbf{y}_n^s \,|\, \mathbf{z}_n^s = i, \tilde{\Theta})\, \mathbb{P}(\mathbf{z}_n^s = i \,|\, \mathbf{z}_n^{s-1}, \mathbf{u}_n^{s-1}, \tilde{\Theta}) \\
&\quad = [\tilde{\mathcal{Q}}]_{i, \mathbf{z}_n^{s+1}, \mathbf{u}_n^s}\, \tilde{\mathcal{G}}_{\mathbf{y}_n^s, \mathbf{u}_n^s}\, \tilde{\mathcal{E}}_{i, \mathbf{y}_n^s}\, [\tilde{\mathcal{Q}}]_{\mathbf{z}_n^{s-1}, i, \mathbf{u}_n^{s-1}}
\end{aligned} \tag{3.39}$$

The boundary equations will be:

$$\begin{aligned}
&\mathbb{P}(\mathbf{z}_n^1 = i \,|\, \mathbf{z}_n^{2:T_n}, \mathbf{u}_n^{1:T_n}, \mathbf{y}_n^{1:T_n}, \tilde{\Theta}) = \mathbb{P}(\mathbf{z}_n^1 = i \,|\, \mathbf{z}_n^2, \mathbf{u}_n^1, \mathbf{y}_n^1, \tilde{\Theta}) \propto \mathbb{P}(\mathbf{z}_n^1 = i,\, \mathbf{z}_n^2, \mathbf{u}_n^1, \mathbf{y}_n^1, \tilde{\Theta}) \\
&\quad = \mathbb{P}(\mathbf{z}_n^2 \,|\, \mathbf{z}_n^1 = i, \mathbf{u}_n^1, \tilde{\Theta})\, \mathbb{P}(\mathbf{u}_n^1 \,|\, \mathbf{y}_n^1, \tilde{\Theta})\, \mathbb{P}(\mathbf{y}_n^1 \,|\, \mathbf{z}_n^1 = i, \tilde{\Theta})\, \mathbb{P}(\mathbf{z}_n^1 = i, \tilde{\Theta}) \\
&\quad = [\tilde{\mathcal{Q}}]_{i, \mathbf{z}_n^2, \mathbf{u}_n^1}\, \tilde{\mathcal{G}}_{\mathbf{y}_n^1, \mathbf{u}_n^1}\, \tilde{\mathcal{E}}_{i, \mathbf{y}_n^1}\, \tilde{\pi}_i
\end{aligned} \tag{3.40}$$

$$\begin{aligned}
&\mathbb{P}(\mathbf{z}_n^{T_n} = i \,|\, \mathbf{z}_n^{1:T_n-1}, \mathbf{u}_n^{1:T_n}, \mathbf{y}_n^{1:T_n}, \tilde{\Theta}) = \mathbb{P}(\mathbf{z}_n^{T_n} = i \,|\, \mathbf{z}_n^{T_n-1}, \mathbf{u}_n^{T_n-1}, \mathbf{u}_n^{T_n}, \mathbf{y}_n^{T_n}, \tilde{\Theta}) \\
&\quad \propto \mathbb{P}(\mathbf{z}_n^{T_n} = i, \mathbf{u}_n^{T_n}, \mathbf{y}_n^{T_n} \,|\, \mathbf{z}_n^{T_n-1}, \mathbf{u}_n^{T_n-1}, \tilde{\Theta}) \\
&\quad = \mathbb{P}(\mathbf{u}_n^{T_n} \,|\, \mathbf{y}_n^{T_n}, \tilde{\Theta})\, \mathbb{P}(\mathbf{y}_n^{T_n} \,|\, \mathbf{z}_n^{T_n} = i, \tilde{\Theta})\, \mathbb{P}(\mathbf{z}_n^{T_n} = i \,|\, \mathbf{z}_n^{T_n-1}, \mathbf{u}_n^{T_n-1}, \tilde{\Theta}) \\
&\quad = [\tilde{\mathcal{Q}}]_{\mathbf{z}_n^{T_n-1}, i, \mathbf{u}_n^{T_n-1}}\, \tilde{\mathcal{G}}_{\mathbf{y}_n^{T_n}, \mathbf{u}_n^{T_n}}\, \tilde{\mathcal{E}}_{i, \mathbf{y}_n^{T_n}}
\end{aligned} \tag{3.41}$$

Notice that we used d-separation of the graph nodes for simplifying the conditional distributions (refer to Figure 3.9). The normalized conditional probabilities can now be calculated using equations 3.39, 3.40, 3.41 by dividing each probability by the sum of all the similar probabilities

for different values of i.

In practice, we need to generate a separate set of samples for any given patient n, which is utilized for approximating $\gamma_n^s(\cdot)$ and $\nu_n^s(\cdot,\cdot)$ for $1 \leq s \leq T_n$. Suppose $m = 1, 2, \cdots$ iterates over the samples that we generate. Furthermore, let $\mathbf{z}_{n,(m)}^s$ be the value of $\mathbf{z}$ at time period s for the m-th sample that we generate. We follow an ordered Gibbs sampling strategy that naturally performs thinning with factor T_n. Specifically, in any iteration m, we sequentially update the values of all the nodes $\mathbf{z}_n^s$ for $s = 1, \cdots, T_n$ one by one, according to the normalized versions of equations 3.39, 3.40, and 3.41. Afterward, we consider the entire series $\mathbf{z}_n^{1:T_n}$ as our m-th sample. In other words, we update all the dimensions from sample m to sample $m+1$, rather than updating only one dimension. We start the Gibbs sampling procedure with a random instance $\mathbf{z}_{n,(0)}^{1:T_n}$ that is generated according to T_n independent uniform distributions over the set $\{1, \cdots, T_n\}$. We throw away the first B samples that we generate in the burn-in period and use the next M samples for approximating the posterior probabilities:

$$\gamma_n^s(i) = \mathbb{E}[\mathbb{I}\{\mathbf{z}_n^s = i\} \mid \mathbf{u}_n^{1:T_n}, \mathbf{y}_n^{1:T_n}, \tilde{\Theta}] \approx \frac{1}{M} \sum_{m=B+1}^{B+M} \mathbb{I}\{\mathbf{z}_{n,(m)}^s = i\} \tag{3.42}$$

$$\nu_n^s(i,k) = \mathbb{E}[\mathbb{I}\{\mathbf{z}_n^s = i, \mathbf{z}_n^{s+1} = k\} \mid \mathbf{u}_n^{1:T_n}, \mathbf{y}_n^{1:T_n}, \tilde{\Theta}] \approx \frac{1}{M} \sum_{m=B+1}^{B+M} \mathbb{I}\{\mathbf{z}_{n,(m)}^s = i, \mathbf{z}_{n,(m)}^{s+1} = k\} \tag{3.43}$$

A detailed description of the parameter estimation procedure is presented in Algorithm 4. We set the number of burn-in samples as $B = 0.1\,M$.

3.11.2 Numerical Experiments

To validate our proposed approach, we generate synthetic data and implement the MCEM algorithm[4]. We then compare the results with the direct approach (i.e. using the forward-backward algorithm) in different cases. Specifically, we generate N IID samples from the graphical model presented in Figure 3.9. For each sample, we simulate the number of visits

[4]Refer to the DT-HMM directory in *https://github.com/saeedghodsi93/Disease_Progression_Modeling_HMM*

Algorithm 4: An MCEM algorithm for learning the DT-HMM parameters

Input: Initial parameters $\tilde{\Theta} = (\tilde{\pi}, \tilde{\mathcal{Q}}, \tilde{\mu}, \tilde{\eta})$
Output: Estimated parameters $\hat{\Theta} = (\hat{\pi}, \hat{\mathcal{Q}}, \hat{\mu}, \hat{\eta})$
Data: A set of N IID samples, each including $(\mathbf{y}_n^s, \mathbf{u}_n^s)\big|_{1 \leq s \leq T_n}$ for $1 \leq n \leq N$
while not converged **do**

E-step (forward-backward): for $1 \leq n \leq N$

- Calculate $\alpha_n^s(i)$ and $\beta_n^s(i)$
- Calculate $\gamma_n^s(i)$ and $\nu_n^s(i,k)$

E-step (Gibbs sampling): for $1 \leq n \leq N$

- Generate $\mathbf{z}_{n,(0)}^{1:T_n}$ following IID distributions $Unif(\{1, \cdots, I\})$
- For $1 \leq m \leq B + M$, generate $\mathbf{z}_{n,(m)}^{1:T_n}$ according to equations 3.39, 3.40, 3.41
- Remove burn-in samples and approximate $\gamma_n^s(\cdot)$, $\nu_n^s(\cdot,\cdot)$ via equations 3.42, 3.43

M-step:

- Update the parameters as $\hat{\pi}_i = \frac{\sum_{n=1}^N \gamma_n^1(i)}{\sum_{i'=1}^I \sum_{n=1}^N \gamma_n^1(i')}$, $\hat{\mu}_i = \frac{\sum_{n=1}^N \sum_{s=1}^{T_n} \mathbf{y}_n^s \gamma_n^s(i)}{(J-1)\sum_{n=1}^N \sum_{s=1}^{T_n} \gamma_n^s(i)}$, $\hat{\eta}_j = \frac{\sum_{n=1}^N \sum_{s=1}^{T_n} \mathbf{u}_n^s \mathbb{I}\{\mathbf{y}_n^s = j\}}{(L-1)\sum_{n=1}^N \sum_{s=1}^{T_n} \mathbb{I}\{\mathbf{y}_n^s = j\}}$, $[\hat{\mathcal{Q}}]_{i,k,l} = \frac{\sum_{n=1}^N \sum_{s=1}^{T_n - 1} \nu_n^s(i,k)\, \mathbb{I}\{\mathbf{u}_n^s = l\}}{\sum_{\tilde{k}=1}^I \sum_{n=1}^N \sum_{s=1}^{T_n - 1} \nu_n^s(i,\tilde{k})\, \mathbb{I}\{\mathbf{u}_n^s = l\}}$

return $\hat{\pi} = \tilde{\pi}$, $\hat{\mathcal{Q}} = \tilde{\mathcal{Q}}$, $\hat{\mu} = \tilde{\mu}$, $\hat{\eta} = \tilde{\eta}$

according to a rounded Normal distribution $T_n \sim round(\mathcal{N}(\mu_T, \sigma_T^2))$. We assume that we have at least two visits per patient. Let $I = 3$, $J = 10$, and $L = 3$. The initial underlying health state node is generated as $\mathbf{z}_n^1 \sim Cat(\pi^*)$. The first physician observation variable will then be generated according to $\mathbf{y}_n^1 \sim Cat(\mathcal{E}^*(\mathbf{z}_n^1, \cdot))$. Similarly, we generate the first intervention variable as $\mathbf{u}_n^1 \sim Cat(\mathcal{G}^*(\mathbf{y}_n^1, \cdot))$. Now, we iterate over the time period index $s = 2, \cdots, T_n$ and generate the underlying health state variable, physician observation variable, and the intervention variable. Specifically, we first generating the health state variable according to $\mathbf{z}_n^s \sim Cat([\tilde{Q}]_{\mathbf{z}_n^{s-1}, \cdot, \mathbf{u}_n^{s-1}})$. Afterward, we generate the physician observation variable as $\mathbf{y}_n^s \sim Cat(\mathcal{E}^*(\mathbf{z}_n^s, \cdot))$ and the intervention variable as $\mathbf{u}_n^s \sim Cat(\mathcal{G}^*(\mathbf{y}_n^s, \cdot))$. Figure 3.10 presents an example of the samples that we've generated following this procedure.

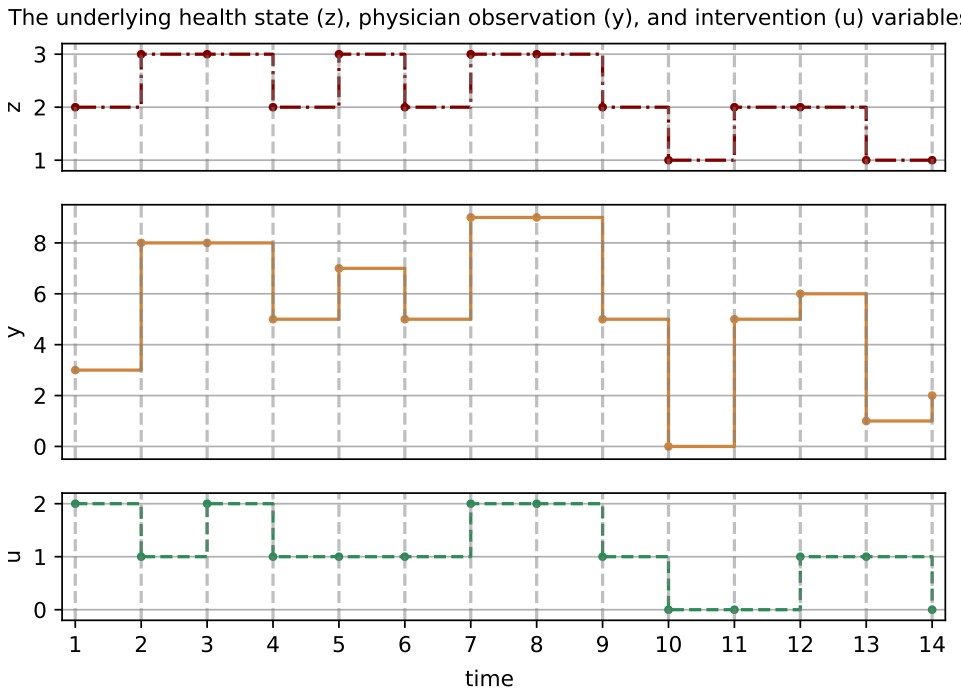

Figure 3.10: An example of the generated DT-HMM synthetic data.

We chose the mean and standard deviations of the number of visits per patient as $\mu_T = 10$, $\sigma_T = 3$. We then chose the following transition probability matrix:

$$\mathcal{Q}^*_{\cdot,\cdot,0} = \begin{bmatrix} 0.20 & 0.45 & 0.35 \\ 0.10 & 0.30 & 0.60 \\ 0.10 & 0.15 & 0.75 \end{bmatrix}, \qquad \mathcal{Q}^*_{\cdot,\cdot,1} = \begin{bmatrix} 0.40 & 0.40 & 0.20 \\ 0.30 & 0.25 & 0.45 \\ 0.10 & 0.30 & 0.60 \end{bmatrix}, \qquad \mathcal{Q}^*_{\cdot,\cdot,2} = \begin{bmatrix} 0.55 & 0.30 & 0.15 \\ 0.40 & 0.35 & 0.25 \\ 0.20 & 0.45 & 0.35 \end{bmatrix}$$

where the first two indices iterate over $i \in \{1, \cdots, I\}$, $k \in \{1, \cdots, I\}$, and the last index iterates over the set of intervention options $l \in \{0, \cdots, L-1\}$. We also set the true initial health state distribution parameter as $\pi^* = (0.11, 0.21, 0.67)$, which is the stationary distribution of $\mathcal{Q}^*_{\cdot,\cdot,0}$. Furthermore, we set the emission probabilities as $\mu^* = (0.15, 0.5, 0.8)$. Ultimately, the intervention probabilities, associated with different observations will be set to:

$$\eta^* = (0.04, 0.15, 0.26, 0.32, 0.43, 0.51, 0.62, 0.77, 0.81, 0.90)$$

As our initial parameters for the EM algorithm, we set $\tilde{\pi} = (0.33, 0.33, 0.33)$. Moreover, we choose the initial transition probability matrix as:

$$\tilde{\mathcal{Q}}_{\cdot,\cdot,0} = \begin{bmatrix} 0.60 & 0.20 & 0.20 \\ 0.40 & 0.40 & 0.20 \\ 0.40 & 0.30 & 0.30 \end{bmatrix}, \quad \tilde{\mathcal{Q}}_{\cdot,\cdot,1} = \begin{bmatrix} 0.20 & 0.50 & 0.30 \\ 0.50 & 0.30 & 0.20 \\ 0.20 & 0.50 & 0.30 \end{bmatrix}, \quad \tilde{\mathcal{Q}}_{\cdot,\cdot,2} = \begin{bmatrix} 0.20 & 0.40 & 0.40 \\ 0.40 & 0.10 & 0.50 \\ 0.30 & 0.10 & 0.60 \end{bmatrix}$$

Finally, set $\tilde{\mu} = (0.45, 0.65, 0.7)$ and $\tilde{\eta} = (0.01, 0.03, 0.06, 0.08, 0.45, 0.48, 0.51, 0.55, 0.91, 0.97)$.

We plot RMSE for different dataset sizes and different number of Monte Carlo samples as a function of the EM iteration in figure 3.11. Although there is an initial jump in the error terms associated with some of the parameters, we observe that all the estimation errors would eventually converge to zero if the dataset size, the number of Monte Carlo samples, and the number of iterations are large enough. Moreover, we can study the effect of the number of Monte Carlo samples in the overall estimation accuracy by comparing the results of Gibbs sampler with the results of the forward-backward algorithm. In this case, it seems that $M = 100$ will give us an accurate result if the sample size is large enough, while we might need more samples for moderate sample sizes. Notice that the scale of the errors vary due to the difference between the scale of the parameters and the difference between the quality of the initial points that we've chosen.

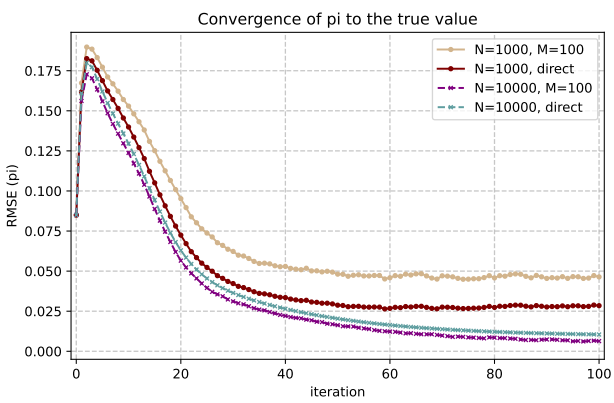

(a) Convergence of $\hat{\pi}$

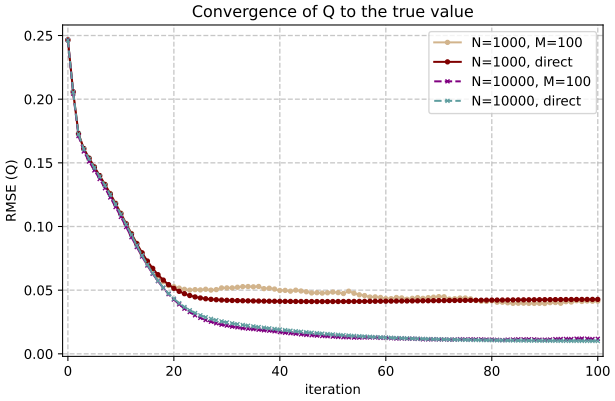

(b) Convergence of $\hat{\mathcal{Q}}$

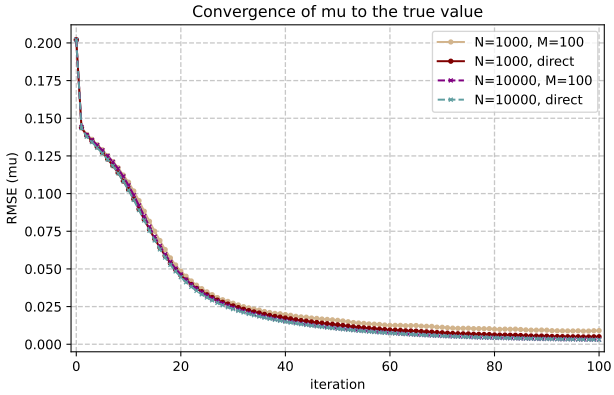

(c) Convergence of $\hat{\mu}$

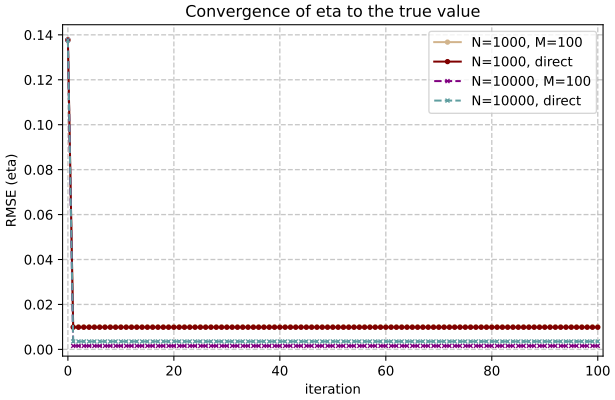

(d) Convergence of $\hat{\eta}$

Figure 3.11: Convergence of the EM estimated parameters to their true values for different sample sizes.

REFERENCES

[AA13] Jamshid Aghaei and Mohammad-Iman Alizadeh. "Demand response in smart electricity grids equipped with renewable energy sources: A review." Renewable and Sustainable Energy Reviews, **18**:64–72, 2013.

[ACW17] Majid Al-Gwaiz, Xiuli Chao, and Owen Q Wu. "Understanding how generation flexibility and renewable energy affect power market competition." Manufacturing & Service Operations Management, **19**(1):114–131, 2017.

[ADM18] Celso Arango, Covadonga M Díaz-Caneja, Patrick D McGorry, Judith Rapoport, Iris E Sommer, Jacob A Vorstman, David McDaid, Oscar Marín, Elena Serrano-Drozdowskyj, Robert Freedman, et al. "Preventive strategies for mental health." The Lancet Psychiatry, **5**(7):591–604, 2018.

[AHS17] Ahmed M Alaa, Scott Hu, and Mihaela Schaar. "Learning from clinical judgments: Semi-markov-modulated marked hawkes processes for risk prognosis." In International Conference on Machine Learning, pp. 60–69. PMLR, 2017.

[AKG12] Hossein Haji Ali Afzali, Jonathan Karnon, and Jodi Gray. "A critical review of model-based economic studies of depression." Pharmacoeconomics, **30**(6):461–482, 2012.

[ALB14] Sabina S Alistar, Elisa F Long, Margaret L Brandeau, and Eduard J Beck. "HIV epidemic control—a model for optimal allocation of prevention and treatment resources." Health care management science, **17**(2):162–181, 2014.

[ALW12] Mahnoosh Alizadeh, Xiao Li, Zhifang Wang, Anna Scaglione, and Ronald Melton. "Demand-side management in the smart grid: Information processing for the power switch." IEEE Signal Processing Magazine, **29**(5):55–67, 2012.

[AMN17] Elodie Adida, Hamed Mamani, and Shima Nassiri. "Bundled payment vs. fee-for-service: Impact of payment scheme on performance." Management Science, **63**(5):1606–1624, 2017.

[AMY10] HA Aalami, M Parsa Moghaddam, and GR Yousefi. "Demand response modeling considering interruptible/curtailable loads and capacity market programs." Applied Energy, **87**(1):243–250, 2010.

[AS19] Ahmed M Alaa and Mihaela van der Schaar. "Attentive State-Space Modeling of Disease Progression." Advances in Neural Information Processing Systems, **32**:11338–11348, 2019.

[ASK20] Bastian Alt, Matthias Schultheis, and Heinz Koeppl. "POMDPs in continuous time and discrete spaces." Advances in Neural Information Processing Systems, **33**:13151–13162, 2020.

[AV18] Ahmed M Alaa and Mihaela Van Der Schaar. "A hidden absorbing semi-markov model for informatively censored temporal data: Learning and inference." The Journal of Machine Learning Research, **19**(1):108–169, 2018.

[BB11] Marie T Brown and Jennifer K Bussell. "Medication adherence: WHO cares?" In Mayo clinic proceedings, volume 86, pp. 304–314. Elsevier, 2011.

[BBD20] János Balogh, József Békési, György Dósa, Leah Epstein, and Asaf Levin. "Online bin packing with cardinality constraints resolved." Journal of Computer and System Sciences, **112**:34–49, 2020.

[BF94] Yoshua Bengio and Paolo Frasconi. "An input output HMM architecture." Advances in neural information processing systems, **7**, 1994.

[BHC16] Philip Burgess, Meredith Harris, Tim Coombs, and Jane Pirkis. "A review of clinician-rated instruments that could be used to assess adults' levels of functioning in specialised public sector mental health services.", 2016.

[BHH82] Gabriel R Bitran, Elizabeth A Haas, and Arnoldo C Hax. "Hierarchical production planning: A two-stage system." Operations Research, **30**(2):232–251, 1982.

[BHP09] Sally C Brailsford, Paul Robert Harper, Brijesh Patel, and Martin Pitt. "An analysis of the academic literature on simulation and modelling in health care." Journal of simulation, **3**(3):130–140, 2009.

[BKT06] Ross Baldick, Sergey Kolos, and Stathis Tompaidis. "Interruptible electricity contracts from an electricity retailer's point of view: valuation and optimal interruption." Operations Research, **54**(4):627–642, 2006.

[BTM12] JW Black, R Tyagi, JS Massey, and JD Williams, U.S. Patent Application No. 12/957,299, 2012.

[CB16] Sarah F Cook and Robert R Bies. "Disease progression modeling: Key concepts and recent developments." Current pharmacology reports, **2**(5):221–230, 2016.

[CBD07] K Cooper, Sally C Brailsford, and Ruth Davies. "Choice of modelling technique for evaluating health care interventions." Journal of the operational research society, **58**(2):168–176, 2007.

[CBS12] J Jaime Caro, Andrew H Briggs, Uwe Siebert, and Karen M Kuntz. "Modeling good research practices—overview: a report of the ISPOR-SMDM Modeling Good Research Practices Task Force–1." Medical Decision Making, **32**(5):667–677, 2012.

[CHK97] Daniel Chisholm, A Healey, and M Knapp. "QALYs and mental health care." Social psychiatry and psychiatric epidemiology, **32**(2):68–75, 1997.

[CKS11] Chen Chen, Shalinee Kishore, and Lawrence V Snyder. "An innovative RTP-based residential power scheduling scheme for smart grids." In 2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 5956–5959. IEEE, 2011.

[CM14] Bibhas Chakraborty and Susan A Murphy. "Dynamic treatment regimes." Annual review of statistics and its application, **1**:447–464, 2014.

[CO16] Clay Campaigne and Shmuel S Oren. "Firming renewable power with demand response: an end-to-end aggregator business model." Journal of Regulatory Economics, **50**(1):1–37, 2016.

[Com16] Michigan Public Service Commission. "Common Demand Response Practices and Program Designs." https://www.michigan.gov/mpsc/commission/workgroups/mi-power-grid/demand-response, 2016.

[Cox72] David R Cox. "Regression models and life-tables." Journal of the Royal Statistical Society: Series B (Methodological), **34**(2):187–202, 1972.

[CW87] Hung-po Chao and Robert Wilson. "Priority service: Pricing, investment, and market organization." The American Economic Review, pp. 899–916, 1987.

[DG12] Meysam Doostizadeh and Hassan Ghasemi. "A day-ahead electricity pricing model based on smart metering and demand-side management." Energy, **46**(1):221–230, 2012.

[DIJ13] Sarang Deo, Seyed Iravani, Tingting Jiang, Karen Smilowitz, and Stephen Samuelson. "Improving health outcomes through better capacity allocation in a community-based chronic care model." Operations Research, **61**(6):1277–1294, 2013.

[DRR15] Sarang Deo, Kumar Rajaram, Sandeep Rath, Uday S Karmarkar, and Matthew B Goetz. "Planning for HIV screening, testing, and care at the Veterans Health Administration." Operations research, **63**(2):287–304, 2015.

[DW11] Benjamin G Druss and E Reisinger Walker. "Mental disorders and medical comorbidity." The Synthesis project. Research synthesis report, pp. 1–26, 2011.

[DYC15] Ruilong Deng, Zaiyue Yang, Mo-Yuen Chow, and Jiming Chen. "A survey on demand response in smart grids: Mathematical models and approaches." IEEE Transactions on Industrial Informatics, **11**(3):570–582, 2015.

[EC08] Oya Eddama and Joanna Coast. "A systematic review of the use of economic evaluation in local decision-making." Health policy, **86**(2-3):129–141, 2008.

[ES12] Olivier Ethgen and Baudouin Standaert. "Population–versus Cohort–Based Modelling Approaches." Pharmacoeconomics, **30**(3):171–181, 2012.

[FDA22] Ali Fattahi, Sriram Dasu, and Reza Ahmadi. "Peak-Load Energy Management by Direct Load Control Contracts." Management Science, 2022.

[FHD20] Joseph Futoma, Michael C Hughes, and Finale Doshi-Velez. "Popcorn: Partially observed prediction constrained reinforcement learning." arXiv preprint arXiv:2001.04032, 2020.

[FHT09] Ahmad Faruqui, Ryan Hledik, and John Tsoukalis. "The power of dynamic pricing." The Electricity Journal, **22**(3):42–56, 2009.

[FKG12] Zhong Fan, Parag Kulkarni, Sedat Gormus, Costas Efthymiou, Georgios Kalogridis, Mahesh Sooriyabandara, Ziming Zhu, Sangarapillai Lambotharan, and Woon Hau Chin. "Smart grid communications: Overview of research challenges, solutions, and standardization activities." IEEE Communications Surveys & Tutorials, **15**(1):21–38, 2012.

[FKS21] Mehdi Fatemi, Taylor W Killian, Jayakumar Subramanian, and Marzyeh Ghassemi. "Medical Dead-ends and Learning to Identify High-risk States and Treatments." Advances in Neural Information Processing Systems, **34**, 2021.

[FMD20] Joseph Futoma, Muhammad A Masood, and Finale Doshi-Velez. "Identifying distinct, effective treatments for acute hypotension with SODA-RL: safely optimized diverse accurate reinforcement learning." AMIA Summits on Translational Science Proceedings, **2020**:181, 2020.

[Fre10] Emma Frew. Applied methods of cost-benefit analysis in health care, volume 4. Oxford University Press, 2010.

[GCW11] Alastair M Gray, Philip M Clarke, Jane L Wolstenholme, and Sarah Wordsworth. Applied methods of cost-effectiveness analysis in healthcare, volume 3. Oxford University Press, 2011.

[Gho22] Saeed Ghodsi. "Efficient Learning of Continuous-Time Hidden Markov Models with Discrete-Time Irregular Observations for Healthcare Intervention Planning.". Master's thesis, University of California, Los Angeles, 2022. Copyright - Database copyright ProQuest LLC.

[GJ79] Michael R Garey and David S Johnson. "A guide to the theory of NP-completeness, A Series of Books in the Mathematical Sciences." Computers and Intractability, WH Freeman and Co, San Francisco, CA, 1979.

[GJM18] Omer Gottesman, Fredrik Johansson, Joshua Meier, Jack Dent, Donghun Lee, Srivatsan Srinivasan, Linying Zhang, Yi Ding, David Wihl, Xuefeng Peng, et al. "Evaluating reinforcement learning algorithms in observational health settings." arXiv preprint arXiv:1805.12298, 2018.

[GLL94] RC Gentleman, JF Lawless, JC Lindsey, and P Yan. "Multi-state Markov models for analysing incomplete disease history data with illustrations for HIV disease." Statistics in medicine, **13**(8):805–821, 1994.

[GS08] Alan M Garber and Jonathan Skinner. "Is American health care uniquely inefficient?" Journal of Economic Perspectives, **22**(4):27–50, 2008.

[GS18] Sherry Glied and Adam Sacarny. "Is the US health care system wasteful and inefficient? A review of the evidence." Journal of Health Politics, Policy and Law, **43**(5):739–765, 2018.

[GTW19] Pengfei Guo, Christopher S Tang, Yulan Wang, and Ming Zhao. "The impact of reimbursement policy on social welfare, revisit rate, and waiting time in a public healthcare system: Fee-for-service versus bundled payment." Manufacturing & Service Operations Management, **21**(1):154–170, 2019.

[HJ05] Asger Hobolth and Jens Ledet Jensen. "Statistical inference in evolutionary models of DNA sequences via the EM algorithm." Statistical applications in genetics and molecular biology, **4**(1), 2005.

[HJ11] Asger Hobolth and Jens Ledet Jensen. "Summary statistics for endpoint-conditioned continuous-time Markov chains." Journal of applied probability, **48**(4):911–924, 2011.

[HLZ19] Ting-Yu Ho, Shan Liu, and Zelda B Zabinsky. "A multi-fidelity rollout algorithm for dynamic resource allocation in population disease management." Health Care Management Science, **22**(4):727–755, 2019.

[HSE16] Haider Tarish Haider, Ong Hang See, and Wilfried Elmenreich. "A review of residential demand response of smart grid." Renewable and Sustainable Energy Reviews, **59**:166–178, 2016.

[Jac11] Christopher H Jackson et al. "Multi-state models for panel data: the msm package for R." Journal of statistical software, **38**(8):1–29, 2011.

[KAA07] Ronald C Kessler, G Paul Amminger, Sergio Aguilar-Gaxiola, Jordi Alonso, Sing Lee, and T Bedirhan Ustun. "Age of onset of mental disorders: a review of recent literature." Current opinion in psychiatry, **20**(4):359, 2007.

[KCD05] Ronald C Kessler, Wai Tat Chiu, Olga Demler, and Ellen E Walters. "Prevalence, severity, and comorbidity of 12-month DSM-IV disorders in the National Comorbidity Survey Replication." Archives of general psychiatry, **62**(6):617–627, 2005.

[KL85] JD Kalbfleisch and Jerald Franklin Lawless. "The analysis of panel data under a Markov assumption." Journal of the American Statistical Association, **80**(392):863–871, 1985.

[Kna03] Martin Knapp. "Hidden costs of mental illness." The British Journal of Psychiatry, **183**(6):477–478, 2003.

[Kna19] Martin Knapp. "Evaluation of Mental Health Interventions." In Oxford Research Encyclopedia of Economics and Finance. Oxford University Press, 2019.

[KO02] Rajnish Kamat and Shmuel S Oren. "Exotic options for interruptible electricity supply contracts." Operations Research, **50**(5):835–850, 2002.

[KOK13] Katarina Kostková, L' Omelina, P Kyčina, and Peter Jamrich. "An introduction to load management." Electric Power Systems Research, **95**:184–191, 2013.

[KVM11] Beste Kucukyazici, Vedat Verter, and Nancy E Mayo. "An analytical framework for designing community-based care for chronic diseases." Production and Operations Management, **20**(3):474–488, 2011.

[LC01] Richard A Levine and George Casella. "Implementations of the Monte Carlo EM algorithm." Journal of Computational and Graphical Statistics, **10**(3):422–439, 2001.

[LDD17] Nancy H Liu, Gail L Daumit, Tarun Dua, Ralph Aquila, Fiona Charlson, Pim Cuijpers, Benjamin Druss, Kenn Dudek, Melvyn Freeman, Chiyo Fujii, et al. "Excess mortality in persons with severe mental disorders: a multilevel intervention framework and priorities for clinical practice, policy and research agendas." World psychiatry, **16**(1):30–40, 2017.

[LDG86] H Stephen Leff, Maqbool Dada, and Stephen C Graves. "An LP planning model for a mental health community support system." Management Science, **32**(2):139–155, 1986.

[LKT20] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. "Offline reinforcement learning: Tutorial, review, and perspectives on open problems." arXiv preprint arXiv:2005.01643, 2020.

[LLL15] Yu-Ying Liu, Shuang Li, Fuxin Li, Le Song, and James M Rehg. "Efficient learning of continuous-time hidden markov models for disease progression." In Advances in neural information processing systems, pp. 3600–3608, 2015.

[LLV19] Elliot Lee, Mariel S Lavieri, and Michael Volk. "Optimal screening for hepatocellular carcinoma: A restless bandit model." Manufacturing & Service Operations Management, **21**(1):198–212, 2019.

[LM13] Jane M Lange and Vladimir N Minin. "Fitting and interpreting continuous-time latent Markov models for panel data." Statistics in medicine, **32**(26):4581–4595, 2013.

[LSS20] MingYu Lu, Zachary Shahn, Daby Sow, Finale Doshi-Velez, and Li-wei H Lehman. "Is Deep Reinforcement Learning Ready for Practical Applications in Healthcare? A Sensitivity Analysis of Duel-DDQN for Hemodynamic Management in Sepsis Patients." In AMIA Annual Symposium Proceedings, volume 2020, p. 773. American Medical Informatics Association, 2020.

[McC95] Paul McCrone. "Predicting mental health service use: Diagnosis based systems and alternatives." Journal of Mental Health, **4**(1):31–40, 1995.

[MCP13] Cathrine Mihalopoulos, ROB Carter, Jane Pirkis, and Theo Vos. "Priority-setting for mental health services." Journal of Mental Health, **22**(2):122–134, 2013.

[Meu14] Benedict De Meulemeester. "Capacity payments: expensive solution for a non-existing problem." EnergyPost. eu, 2014.

[MHS07] Philipp Metzner, Illia Horenko, and Christof Schütte. "Generator estimation of Markov jump processes based on incomplete observations nonequidistant in time." Physical Review E, **76**(6):066702, 2007.

[Mou12] DR Mould. "Models for disease progression: new approaches and uses." Clinical Pharmacology & Therapeutics, **92**(1):125–131, 2012.

[MVP11] Cathrine Mihalopoulos, Theo Vos, Jane Pirkis, and Rob Carter. "The economic analysis of prevention in mental health programs." Annual review of clinical psychology, **7**:169–201, 2011.

[NL92] Sang-jin Nam and Rasaratnam Logendran. "Aggregate production planning—a survey of models and methodologies." European Journal of Operational Research, **61**(3):255–272, 1992.

[NS17] Karthik V Natarajan and Jayashankar M Swaminathan. "Multi-treatment inventory allocation in humanitarian health settings under funding constraints." Production and Operations Management, **26**(6):1015–1034, 2017.

[NSP12] Samuel Newell, Kathleen Spees, Johannes Pfeifenberger, Robert Mudge, Michael DeLucia, and Robert Carlton. "ERCOT investment incentives and resource adequacy." The Brattle Group, prepared for the Electric Reliability Council of Texas, 2012.

[Org06] World Health Organization. "Economic aspects of the mental health system: key messages to health planners and policy-makers." Technical report, World Health Organization, 2006.

[PB15] Heikki Peura and Derek W Bunn. "Dynamic pricing of peak production." Operations Research, **63**(6):1262–1279, 2015.

[PBF13] Alex Papalexopoulos, Jacob Beal, and Steven Florek. "Precise mass-market energy demand management through stochastic distributed computing." IEEE Transactions on Smart Grid, **4**(4):2017–2027, 2013.

[PD11] Peter Palensky and Dietmar Dietrich. "Demand side management: Demand response, intelligent energy systems, and smart loads." IEEE transactions on industrial informatics, **7**(3):381–388, 2011.

[PFU18] Ryan Palmer, Naomi J Fulop, and Martin Utley. "A systematic literature review of operational research methods for modelling patient flow and outcomes within community healthcare and other settings." Health Systems, **7**(1):29–50, 2018.

[PGT03] Joelle Pineau, Geoff Gordon, Sebastian Thrun, et al. "Point-based value iteration: An anytime algorithm for POMDPs." In IJCAI, volume 3, pp. 1025–1032. Citeseer, 2003.

[PK16] Michael E Porter, Robert S Kaplan, et al. "How to pay for health care." Harv Bus Rev, **94**(7-8):88–98, 2016.

[Pos15] Barry Posner. "The Fundamentals of Electricity Markets." https://www.e-education.psu.edu/ebf200/node/151, 2015.

[PPH16] Daryl Pritchard, Allison Petrilla, Shawn Hallinan, Donald H Taylor Jr, Vernon F Schabert, and Robert W Dubois. "What contributes most to high health care costs? health care spending in high resource patients." Journal of Managed Care & Specialty Pharmacy, **22**(2):102–109, 2016.

[Qdr06] QJUDE Qdr. "Benefits of demand response in electricity markets and recommendations for achieving them." US Dept. Energy, Washington, DC, USA, Tech. Rep, **2006**, 2006.

[RBK05] Renée Romeo, Sarah Byford, and Martin Knapp. "Annotation: Economic evaluations of child and adolescent mental health interventions: a systematic review." Journal of Child Psychology and Psychiatry, **46**(9):919–930, 2005.

[RCO09] Nerea Ruiz, Iñigo Cobelo, and José Oyarzabal. "A direct load control model for virtual power plant management." IEEE Transactions on Power Systems, **24**(2):959–966, 2009.

[RV08] Badri Ramanathan and Vijay Vittal. "A framework for evaluation of advanced direct load control with minimum disruption." IEEE Transactions on Power Systems, **23**(4):1681–1688, 2008.

[Sag18] Soroush Saghafian. "Ambiguous partially observable Markov decision processes: Structural results and applications." Journal of Economic Theory, **178**:1–35, 2018.

[SAM14] SAMHSA. "Projections of national expenditures for treatment of mental and substance use disorders, 2010-2020." Substance Abuse and Mental Health Services Administration, 2014.

[SAM18] SAMHSA. "Key substance use and mental health indicators in the United States: Results from the 2017 National Survey on Drug Use and Health." Substance Abuse and Mental Health Services Administration, 2018.

[SB19] Nur Sunar and John R Birge. "Strategic commitment to a production schedule with uncertain supply and demand: Renewable energy in day-ahead electricity markets." Management Science, **65**(2):714–734, 2019.

[SCS14] Lachlan Standfield, Tracy Comans, and Paul Scuffham. "Markov modeling and discrete event simulation in health care: a systematic comparison." International journal of technology assessment in health care, **30**(2):165, 2014.

[SCS20] Kristen A Severson, Lana M Chahine, Luba Smolensky, Kenney Ng, Jianying Hu, and Soumya Ghosh. "Personalized Input-Output Hidden Markov Models for Disease Progression Modeling." In Machine Learning for Healthcare Conference, pp. 309–330. PMLR, 2020.

[sei16] "Solar Energy Industries Association." https://www.seia.org/, 2016.

[SGL19] Zhaonan Sun, Soumya Ghosh, Ying Li, Yu Cheng, Amrita Mohan, Cristina Sampaio, and Jianying Hu. "A probabilistic disease progression modeling approach and its application to integrated Huntington's disease observational data." JAMIA open, **2**(1):123–130, 2019.

[Sia14] Pierluigi Siano. "Demand response and smart grids—A survey." Renewable and sustainable energy reviews, **30**:461–478, 2014.

[Sie68] DAVID Siegmund. "On the asymptotic normality of one-sided stopping rules." The Annals of Mathematical Statistics, **39**(5):1493–1497, 1968.

[SJ13] Steven P Segal and Leah A Jacobs. "Deinstitutionalization." In Encyclopedia of social work. Oxford University Press, 2013.

[SLL11] Susan M Shortreed, Eric Laber, Daniel J Lizotte, T Scott Stroup, Joelle Pineau, and Susan A Murphy. "Informing sequential clinical decision-making through reinforcement learning: an empirical study." Machine learning, **84**(1-2):109–136, 2011.

[SNS18] Anand Satiani, Julie Niedermier, Bhagwan Satiani, and Dale P Svendsen. "Projected workforce of psychiatrists in the United States: A population analysis." Psychiatric Services, **69**(6):710–713, 2018.

[Son78] Edward J Sondik. "The optimal control of partially observable Markov processes over the infinite horizon: Discounted costs." Operations research, **26**(2):282–304, 1978.

[SPK13] Guy Shani, Joelle Pineau, and Robert Kaplow. "A survey of point-based POMDP solvers." Autonomous Agents and Multi-Agent Systems, **27**(1):1–51, 2013.

[SS21] Nur Sunar and Jayashankar M Swaminathan. "Net-metered distributed renewable energy: A peril for utilities?" Management Science, **67**(11):6716–6733, 2021.

[SSS16] Maryam H Shoreh, Pierluigi Siano, Miadreza Shafie-khah, Vincenzo Loia, and João PS Catalão. "A survey of industrial applications of Demand Response." Electric Power Systems Research, **141**:31–49, 2016.

[STY19] Nicos Savva, Tolga Tezcan, and Özlem Yıldız. "Can yardstick competition reduce waiting times?" Management Science, **65**(7):3196–3215, 2019.

[TB11] R Tyagi and JW Black, U.S. Patent Application No. 12/790,655, 2011.

[TEG15] E Fuller Torrey, K Entsminger, J Geller, J Stanley, and DJ Jaffe. "The shortage of public hospital beds for mentally ill persons." Montana, **303**(20.9):6–9, 2015.

[Tem18] James Temple. "The $2.5 trillion reason we can't rely on batteries to clean up the grid." Retrieved January, **17**:2020, 2018.

[Tor06] George W Torrance. "Utility measurement in healthcare." Pharmacoeconomics, **24**(11):1069–1078, 2006.

[USS03] Edwina S Uehara, Debra Srebnik, and Michael Smukler. "Statistical and consensus-based strategies for grouping consumers in mental health level-of-care schemes." Administration and Policy in Mental Health and Mental Health Services Research, **30**(4):287–306, 2003.

[Van78] Charles Van Loan. "Computing integrals involving the matrix exponential." IEEE transactions on automatic control, **23**(3):395–404, 1978.

[VZV14] John S Vardakas, Nizar Zorba, and Christos V Verikoukis. "A survey on demand response programs in smart grids: Pricing methods and optimization algorithms." IEEE Communications Surveys & Tutorials, **17**(1):152–178, 2014.

[Wis11] Matthias Wissner. "The Smart Grid–A saucerful of secrets?" Applied Energy, **88**(7):2509–2518, 2011.

[WS19] Erwin Walraven and Matthijs TJ Spaan. "Point-based value iteration for finite-horizon POMDPs." Journal of Artificial Intelligence Research, **65**:307–341, 2019.

[WSW14] Xiang Wang, David Sontag, and Fei Wang. "Unsupervised learning of disease progression models." In Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 85–94, 2014.

[WWG10] Qiuwei Wu, Peng Wang, and Lalit Goel. "Direct load control (DLC) considering nodal interrupted energy assessment rate (NIEAR) in restructured power systems." IEEE Transactions on Power Systems, **25**(3):1449–1456, 2010.

[YLN21] Chao Yu, Jiming Liu, Shamim Nemati, and Guosheng Yin. "Reinforcement learning in healthcare: A survey." ACM Computing Surveys (CSUR), **55**(1):1–36, 2021.