Personalized Medicine in Chronic Disease Managment

by

Greggory J. Schell

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial and Operations Engineering)
in The University of Michigan
2015

Doctoral Committee:

Assistant Professor Mariel S. Lavieri, Chair
Associate Professor Amy Cohn
Professor Rodney A. Hayward
Professor Susan A. Murphy
Professor Mark P. Van Oyen

ACKNOWLEDGEMENTS

Firstly, I would like to acknowledge the contributions of all of my collaborators, whose invaluable input and expertise helped make this dissertation possible. I would like to thank Dr. Joshua D. Stein for his domain knowledge and efforts toward the research on glaucoma, especially his work in writing the discussion section on the monitoring algorithm. Thank you to Prof. Jonathan E. Helm for your work on developing the Kalman filter. I would also like to thank Drs. Jeremy B. Sussman and Rodney A. Hayward for their openness to working with an eager engineering student, like myself, to model and solve cardiovascular disease problems. Their expertise in the field and their commitment to improving medical decision making in clinical practice is truly inspirational. I would also like to thank Dr. Wyndy Wiitala for her help with collecting and accessing the VA database. And of course, I want to thank my advisor, Prof. Mariel S. Lavieri, whose guidance, knowledge and ingenuity has been critical to my success as a student and will continue to shape my academic future. I also want to thank everyone, especially my committee members, who reviewed and provided feedback for this dissertation.

I would also like to thank the lifelong friends I made while a student at Michigan. Thank you to Zohar Strinka, my officemate and confidant, who has been a great friend and colleague over the past five years. Thanks Troy Long for your amazing friendship and for hosting game nights! I also want to thank Jivan Deglise-Favre-Hawkinson for being a close friend and excellent conference partner. And thank you

to every other friend I've made along the way.

Finally, I want to thank my family, especially my parents Kimberly Schell and David Schell, for pushing me to excel throughout my career. As the first in my immediate family to graduate with a bachelor's degree, I am ever-so appreciative of my parent's unwavering desire to see me succeed academically and professionally. And I want to thank my loving boyfriend, Zachary Blackwell, for supporting me each and every day.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

CHAPTER I

Introduction

Chronic diseases are persistent medical conditions which affect half of all adults in the United States [151]. Chronic diseases are responsible for the majority of deaths [107] and 84% of all health care spending [7] in the United States. The U.S. National Center for Health Statistics defines a chronic disease as a disease which lasts between 3 months and the patient's lifetime, which typically cannot be cured by medication [51]. Some chronic diseases, like glaucoma, are severely debilitating and diminish the quality of life of the patient [54], while others such as cardiovascular disease are both debilitating and a leading cause of death in the United States [77].

The nature of these long-term chronic conditions present monitoring and treatment challenges to practicing clinicians and medical researchers. As each patient is different, a key challenge is how to use information learned about each patient's disease characteristics over time to tailor monitoring and treatment decisions. Medical care for chronic disease requires many sequential decisions where each decision has strong implications for future decisions. Furthermore, long-term clinical management also leads to decreasing rates of adherence to medical interventions.

The medical community has begun to pursue personalized medicine, i.e. tailoring clinical decisions to the patient's individual disease characteristics, as a way to im-

prove health outcomes for patients with chronic disease [61], [28], [112], [48], [156]. Working together, operations researchers and clinicians can address the challenges of personalized medicine and chronic care. Noise extraction and filtering methods can be used to improve the understanding of disease dynamics. Methods, such as dynamic programming, enable the tractable modeling and optimization of sequential decision making problems. Operations research can be used to tailor treatment plans to the level of adherence and optimize how incentives are allocated to improve their adherence. Finally, with increasing health care spending, operations researchers can develop models which are sensitive to resource limitations at both the patient and population level.

By combining operations research with the principles of personalized medicine, we have developed novel mathematical models to answer high impact clinical questions faced when managing patients with chronic conditions. This has lead to the creation of decision support tools for real world implementation. We began our research by understanding how information about a single patient can be used to personalize the patient's forecasted disease dynamics and likelihood of disease progression. Next, we considered how models of heterogeneity in disease characteristics and patient behavior can be embedded within an optimization framework to design individualized treatment plans. Finally, we developed a resource allocation model for improving patient adherence to their individualized treatment plans. Figure 1.1 summarizes the three chapters of this dissertation.

1.1 Disease Progression Identification for Patients with Glaucoma

Chapter II addresses the problem of utilizing correlated repeated measures data affected by measurement and process noise in statistical classification models. By

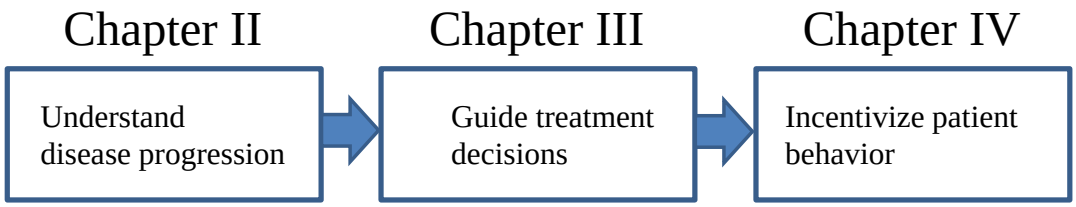

Figure 1.1: Overview of dissertation chapters.

extracting noise from the data prior to model construction, we can reduce misclassification rates when identifying patients with worsening diseases and declining health. This improved classification model was then embedded within a dynamic and personalized monitoring algorithm for determining the time between monitoring visits for chronic disease. The work has been applied to glaucoma patients using clinical trial data for training and testing the models. We found that our classification model improved the area under the receiver operating characteristic curve from 0.889 (using traditional methods) to 0.961. The increased model accuracy then translated to improvements for the monitoring algorithm as compared to current clinical guidelines for glaucoma patients: 29% increase in efficiency of testing and 57% reduction in diagnostic detection of disease progression.

1.2 Treatment Planning for Cardiovascular Disease

Adherence to pharmacotherapy is a major concern for both health care and insurance providers. Poor adherence to medication can lead to avoidable adverse health events, such as heart attack and stroke. In Chapter III, we developed a dynamic programming formulation for incorporating patient adherence to medication, along with individualized risk factors, when deriving optimal treatment policies. Our ap-

proach utilized a coherent risk measure in the calculation of rewards to control the marginal benefit of adherence on treatment outcomes. This work has been applied to hypertension treatment planning for patients at risk of cardiovascular disease. We first considered treatment policies which minimized a patient's expected number of coronary heart disease events using observational data from the U.S Department of Veterans Affairs. Our personalized and optimal treatment decisions lead to a 9.4% reduction in the expected number of coronary heart disease events when compared to current treatment guidelines. Next, we considered treatment policies which maximized a patient's expected quality-adjusted life years (QALYs) using data from a nationally-representative survey, NHANES. We also analyzed the impact of physician knowledge of the treatment disutility incurred by a patient when taking medication. We found that misestimation of disutility can lead to upwards of 26.9% of patients being mistreated for hypertension. Lastly, we studied the effects of time-varying treatment benefit and disutility on current and future risk treatment strategies. Our analysis indicated that the future risk strategy based on a forecasted 5% 10-year cardvioascular risk is the most sensitive to changes in how long a medication must be taken before maximum treatment benefit is attained. This phenomenon is due to the lengthier expected treatment duration for patients under the future 5% risk strategy.

1.3 Copayment Restructuring for Improved Adherence to Treatment Plans

Building upon the theoretical and numerical results of the dynamic programming formulation for hypertension treatment planning, we developed a bilevel optimization model for determining allocations of resources designed to improve patient adherence to their individualized treatment plan. This model considers allocating resources

across a population which is heterogeneous in their risk factors for cardivascular disease. The bilevel model studies the role of budgetary and equitability constraints on the optimal allocations. Chapter IV develops an iterative nonlinear programming technique to solve this resource allocation problem. Our research indicates that targeting resources based on individual characteristics can improve health outcomes (4.35 quality-adjusted life years per 1000 patients) for the patient population and reduce total health costs.

1.4 Conclusion and Future Research

This research has developed mathematical methods for addressing key issues in the management of chronic disease patients, including identifying disease progression, determining monitoring frequency, developing treatment plans, and allocating incentives for improving adherence to medication. This body of work has generated advances in engineering and medical decision making with many further opportunities for continued research. Chapter V summarizes the contributions of the work and presents the details of some proposed future research.

CHAPTER II

Disease Progression Identification for Patients with Glaucoma

In this chapter, we present a statistical classification algorithm for noisy and correlated data, and we highlight how this algorithm may be used to guide decision making. The algorithm assesses the probability that a patient will experience disease progression by combining a Kalman filter approach (to reduce noise and model the system dynamics) with generalized estimating equations (to account for the correlation of repeated observations for a patient). We show that, by optimally extracting noise in the data prior to parameterizing our algorithm, it is possible to reduce misclassification error rates. This work has been published in *BMC Medical Informatics and Decision Making* [138]. Next, we utilize the algorithm to assist clinicians in developing monitoring schedules for their patients. Our approach yields a decision support system for monitoring glaucoma patients with noisy and correlated observations. This extension has been published in *Ophthalmology* [137]. While our work is directly motivated by open angle glaucoma, the algorithm and decision support system are generalizable to chronic diseases where process and measurement noise affect longitudinal data collection and interpretation.

2.1 Background

Open angle glaucoma (OAG) is a chronic degenerative ocular disease characterized by damage to the optic nerve. It is the second leading cause of blindness with an estimated 2.2 million adult Americans diagnosed with this disease [54], [122]. Patients with OAG are monitored regularly via visual field (VF) tests and intraocular pressure (IOP) readings [90], [109], [17], [46], [106], [159]. Clinicians monitor glaucoma patients in order to determine whether significant disease progression has occurred which would warrant changes in the medical management of the disease (e.g. perform surgery or change medications). VF tests summarize vision loss through perimetric outputs (measurements of light sensitivity in different parts of the eye) such as mean deviation (MD) and pattern standard deviation (PSD). These perimetric outputs compare the mean (standard deviation, respectively) light sensitivity of a given patient's eye to a normative database of patients without glaucoma. IOP readings measure the fluid pressure on the eye caused by aqueous humor, a natural fluid produced to provide nutrition and immune defense to the eye.

Clinicians use past and present VF and IOP tests to detect disease progression. However, the clinical observations are subject to both process and measurement noise. Errors in machine calibration, patient anxiety, human error in administering tests, and variations in measurement technique can all contribute to measurement noise when assessing chronic diseases [91]. Biological variability, like intraday fluctuation of intraocular pressure, is a contributing factor to process noise which can affect the ability to identify significant disease progression [101]. Distinguishing between signal and noise becomes paramount when these noisy measurements are used in decision making.

While most clinicians (glaucoma specialists in particular) are aware of the variability in VF and IOP measurements, non-specialists may not fully appreciate the importance of considering noise in VF findings and IOPs from visit to visit: they may erroneously conclude a patient is progressing or non progressing when they observe variability and make treatment decisions accordingly. Studies have shown that the difficulties associated with evaluating patients with glaucoma to assess for disease progression have led to undertreatment [53],[99] and that decision aids, such as risk calculators [98], are useful supplements to clinician judgment. Our proposed methodology aids clinicians (both specialists and nonspecialists) in their decision making process by providing them with a mathematical framework that systematically accounts for noise in the data used to identify disease progression.

2.2 Literature Review

Statistical classification, such as detecting disease progression, is not new to the statistics and machine learning literature. Some traditional methods for supervised classification include support vector machines [23], random forests [154], and logistic regression [93], [118]. These traditional methods assume independent observations as inputs, whereas identifying significant disease progression requires analysis of correlated longitudinal data of a particular patient. We propose the use of generalized estimating equations (GEE) to statistically model the relationship between OAG health metrics and progression. GEE has been extensively used in the medical literature to: assess improvements from conversion to electronic health records [130], identify risk factors for chronic obstructive pulmonary disease [142], identify predictors of influenza vaccine acceptance [47], and study spatially correlated binary data in neuroimaging [6]. However, applying GEE to raw measurements may lead to deci-

sion making that is informed by noisy observations, not measurements which reflect true disease dynamics.

A key feature of the algorithm developed relies on its application of Kalman filtering to model the disease dynamics. Kalman filtering is a technique for identifying signal in the presence of measurement and process noise [73]. Other modeling techniques, such as Markov chains [132], [65], have been used to model disease dynamics; however, such techniques do not extract noise to estimate the true value of the disease variables. The Kalman filter approach has been used to estimate pulmonary blood flow [22], track cardiovascular signals [105], continuously monitor glucose levels [85], and monitor prostate specific antigen levels in prostate cancer patients [87]. These applications of the Kalman filter are used for predicting important health metrics, but the relationship between filtered estimates of these metrics and significant disease progression identification has not yet been modeled. Hence, in our statistical classification algorithm, we combine GEE for logisitic regression with Kalman filtered estimates to both reduce noise and account for correlation.

The key contribution of this work is the utilization of Kalman filter estimates of patient health metrics in logistic regression models. We show that our classification algorithm improves significant disease progression identification compared to logistic regression models constructed using raw clinical observations. Furthermore, Kalman filter estimates explicitly account for process and measurement noise inherent in clinical data, making these estimates more informative than raw observations alone. This is of particular interest to clinicians who must decide when to monitor patients and select treatments based on the patient's likelihood of progression. While our initial application is to patients with OAG, this methodology is applicable to other chronic diseases.

2.3 Methods

Our proposed methodology combines an understanding of system dynamics properties through a Kalman filter approach with the marginal response technique of GEE to estimate the true value of clinical observations and to improve the ability of the logistic regression models to identify significant OAG progression. Our methodology is as follows: first, we build our dataset for analysis from clinical observations of a randomized clinical trial of OAG patients. Next, we construct a robust definition of significant OAG progression based on the knowledge of subject matter experts. We then apply Kalman filtering to the repeated measures data of the large-scale randomized controlled clinical trial to estimate the true values of variables believed to be correlated to significant OAG progression. The GEE with a logit link function is then applied to the filtered data to develop a probability function for significant OAG progression. Finally, through cross validation, we measure sensitivity and specificity and calculate the area under the receiver operating characteristic (ROC) curve (AUC) to evaluate the accuracy of the probability function at identifying significant OAG progression.

Analyses were run using R version 2.12.2 and MATLAB version 7.7. For all analyses, $p < 0.05$ was considered statistically significant.

2.3.1 Data

The data set used for parameterization and validation of our proposed methodology came from the Collaborative Initial Glaucoma Treatment Study (CIGTS). CIGTS provided clinical visit data for patients with early to moderate OAG over 10 years. All patients were seen approximately every 6 months following initial intervention to have a VF test and IOP check. Longitudinal MD and PSD values from

VF tests, along with longitudinal IOP measurements, were obtained for each patient from the clinical trial data. From the longitudinal measurements, we calculated velocities and accelerations for MD, PSD, and IOP for every patient at each visit. We also extracted demographic information (age, sex, and race) for every patient in the clinical trial.

Our inclusion criteria required patients to have at least 4 follow-up VF tests and IOP checks after initial intervention. We also required the patient's VF test data to include information on each individual light sensitivity point in order to apply our progression definition. Further, patients were required to have been initially treated with medical therapy. Patients were censored when they left the clinical trial or if they underwent trabeculectomy or argon laser trabeculoplasty (ALT).

2.3.2 Progression Labeling

Our modeling approach used repeated measures data from a randomized clinical trial to obtain information about the evolution of OAG over time. Given the set of longitudinal data obtained sequentially for each patient over the course of the clinical trial, we used the input of subject matter experts to retrospectively identify patients that experienced a significant change in disease characteristics that would warrant clinical intervention. It is important to note that our definition of significant disease progression is meant to serve as an alert to clinicians. Clinicians use the information obtained from the models along with their experience and patient-specific factors to ultimately decree how to best care for a particular OAG patient.

A large drop in MD is generally accepted as an instance of OAG progression. Furthermore, we used the Hodapp-Anderson-Parish (HAP) classification [69] in our definition of significant OAG progression. Patients were labeled as experiencing progression at visit j when there was a loss of ≥ 3 decibels (dB) of mean deviation

(MD) with respect to baseline MD at visit j and this loss of MD also occured for some future visit $k : k > j$ or if the patient shifted upward in HAP class (e.g. moderate to severe). We applied this definition to all patients in our dataset. This definition of progression requires both significant change in disease characteristics at the particular visit and a validation of this change in some future visit. Validation of the loss of MD at a future visit mitigates the chance of erroneously concluding a patient is progressing or not progressing due to noise in the data. In practice, this added level of validation necessitates the development of an OAG disease progression probability function, since knowledge about future visits is not available to clinicians when identifying whether a patient has progressed.

Our definition of OAG progression (using either a validated 3 dB decline in MD or worsening based on HAP criteria) was compared against other suggested definitions of OAG progression (validated decline of 3 dB in MD alone [108], progression based on HAP criteria alone [69], and a point-wise linear regression method for progression detection [113]) and the model performed well irrespective of the OAG progression definition chosen. Given our interest in identifying global worsening of visual field from glaucoma as well as segmental areas of the visual field loss from glaucoma we opted to use the more encompassing progression definition, characterized as a validated decline in MD of 3 dB from baseline or worsening based on HAP criteria for our analyses.

2.3.3 Application of Kalman Filter

The raw measurements obtained from sequential testing of OAG patients can be susceptible to process and measurement noise. To mitigate the effect of process and measurement noise, we utilized a Kalman filter approach to estimate the true value for observations obtained from VF and IOP tests. The Kalman filter utilizes

recursive mathematical equations to optimally estimate the mean and covariance parameters of a process to characterize the state of the disease system [73]. The disease state can be multidimensional. We considered our state, $\alpha_{i,j} \in \Re^n$, to be the value of OAG-related variables (MD, PSD, IOP, and their respective velocities and accelerations) at the j^{th} observation time for patient i. The Kalman filter assumes the state is a Gaussian random variable and that the evolution of the disease state is governed by a linear stochastic difference equation:

$$\alpha_{i,j} = \mathbf{T}\alpha_{i,j-1} + w_{i,j-1} \tag{2.1}$$

Equation (2.1) represents the disease state, $\alpha_{i,j}$, of patient i at the current period j, as a transformation of the state of the last period, $\alpha_{i,j-1}$, according to the transition matrix $\mathbf{T}$, plus Gaussian process noise, $w_{i,j-1}$, with mean 0. Additionally, the covariance of the disease state variables for the current period, $\Sigma_{i,j}$, is also governed by the transition matrix $\mathbf{T}$, plus the covariance matrix of the process noise, $\mathbf{Q}$.

$$\Sigma_{i,j} = \mathbf{T}\Sigma_{i,j-1}\mathbf{T}' + \mathbf{Q} \tag{2.2}$$

Moreover, the true state cannot be directly measured. Instead, the clinical observations, $z_{i,j}$, are assumed to be a linear function of the state, $\alpha_{i,j}$, transformed by the matrix $\mathbf{H}$ plus Gaussian measurement noise, $v_{i,j}$, with mean 0:

$$z_{i,j} = \mathbf{H}\alpha_{i,j} + v_{i,j}$$

We obtained state estimates, $\hat{\alpha}_{i,j}$ of the true state $\alpha_{i,j}$ for each patient's visit by recursively predicting the disease state values using a population-based understanding of OAG mechanics (i.e. parameterized transition and covariance matrices) and updating the estimates to reflect the patient's particular disease evolution [73]. First, equation (1) was used to forecast the patient's disease state in the next period. The

Kalman filter then used the patient's observed disease state in order to update the estimate of the true state,$\hat{\alpha}_{i,j}$. This personalized trajectory process was repeated for each patient, over the course of the patient's duration in the clinical trial, to obtain the best estimates of the patient's state at each observation period. This procedure accounted for the inherent process and measurement noise to provide information to the decision makers that better reflects the actual disease state of each patient.

The final Kalman filter was developed using an expectation maximization (EM) algorithm for parameter estimation [57].

2.3.4 Generalized Estimating Equations

We next used GEE for logistic regression to develop the probability function for significant OAG progression. It is an extension of generalized linear models to repeated measures data analysis using quasi-likelihood estimation [161]. GEE is a semiparametric regression technique that uses an iterative algorithm, Newton-Rhapsody, to estimate the coefficient parameters. Unlike linear mixed effects models, GEE is robust to the specification of the correlation structure and requires only the correct specification of the marginal means to obtain consistent and asymptotically normal parameter estimates [60].

Let $y_{i,j}$ be the response (i.e. progression label) for patient i at time j, and let $\mu_{i,j}$ be the expected value of $y_{i,j}$. The GEE approach assumes independence at the patient-level and relates the marginal response, $\mu_{i,j}$, to a linear combination of the covariates, $x_{i,j}$, by the link function, $g(\cdot)$.

$$g(\mu_{i,j}) = \mathbf{x}_{i,j}^T \beta \tag{2.3}$$

where β is a $p \times 1$ vector of unknown regression coefficients.

Traditionally, the covariates used in GEE are obtained from the raw observed

measurements $z_{i,j}$. Our proposed methodology called for using the state estimates $\hat{\alpha}_{i,j}$ as the input to our model. We compared the performance of the GEE model which uses $\hat{\alpha}_{i,j}$ (Kalman filter estimates) against the GEE model which uses $z_{i,j}$ (raw observations).

Through the GEE framework, we considered the variance of the response variable, $V(y_{i,j})$, to be a function, $v(\cdot)$, of the mean response, $\mu_{i,j}$:

$$V(y_{i,j}) = v(\mu_{i,j})\phi \tag{2.4}$$

where $v(\cdot)$ is a known variance function and ϕ is a possibly unknown scale parameter.

Because of our Bernoulli response variable, i.e. 1 for progression and 0 for nonprogression, we used the logit link function, logit variance function, and scale parameter, $\phi = 1$:

$$g(\mu_{i,j}) = log[\frac{\mu_{i,j}}{1-\mu_{i,j}}] \tag{2.5}$$

$$v(\mu_{i,j}) = \mu_{i,j}(1-\mu_{i,j}) \tag{2.6}$$

Repeated measures data are inherently correlated, with independence at the patient level. GEE uses a $n \times n$ "working" correlation matrix, $\mathbf{Z}_i$, for each patient's sequence of response variables $\mathbf{y}_i$, to account for this inherent correlation. We utilized an autoregressive correlation structure. The autoregressive correlation structure assumes a first-order relationship between the measurements. The correlation depends on the magnitude of the time difference between the measurements:

$$Corr(y_{i,s}, y_{i,t}) = \rho^{|s-t|}, \rho \in (-1, 1) \tag{2.7}$$

The final covariate set for the logistic regression was obtained via forward variable selection. Variable selection was initialized with MD, PSD, and change in MD. Chi-

squared tests were used to evaluate the benefit of adding a single variable to the model. We iteratively added the variable with the smallest Chi-square test p-value to the model until no new variables were statistically significant ($\alpha = 0.10$).

2.3.5 Model Performance

To assess the performance of the logistic regression models, we developed receiver operating characteristic (ROC) curves. 10-fold cross validation was performed to calculate sensitivity and specificities at various discrimination thresholds. Receiver operating characteristic (ROC) curves were created using the average sensitivities and specifities across the 10-fold cross validation to compare the performance of the logistic regression model with raw observations as input versus the model with Kalman filter state estimates as input. Estimates of the area under the ROC curve (AUC) were obtained for each iteration of the 10-fold cross validation.

2.4 Results

From the CIGTS data, 90 patients satisfied our inclusion criteria. The mean (standard deviation) number of visits for patients who met the inclusion criteria was 15.1 (2.6) visits. Nearly 99% of the patients had at least 8 follow-up visits.

We calculated the overall patient mean (and standard deviation) of every variable (Kalman filter estimates and raw observations) from the VF and IOP tests (Table 2.1) for instances of progression and nonprogression separately. We note that generally the difference between the progressing and nonprogressing means for a variable is larger for the Kalman filter estimates than in the raw measurement data set, e.g. the difference between progressing and nonprogressing mean PSD is 7.643 for the Kalman filter estimates and 4.831 for the raw observations. The increased difference between progressing and nonprogressing means for a variable in the Kalman filter

data set is due to the linear system dynamics framework of the Kalman filter. As time increases, the linear trajectory of the Kalman filter results in more disparate variable values between nonprogressing and progressing instances compared to the "noisy" trajectories in the raw observations data set.

	Kalman Filter Estimates		**Raw Observations**	
Variable	Progressing	Nonprogressing	Progressing	Nonprogressing
MD	-8.75 (7.38)	1.24 (4.79)	-8.55 (4.30)	-2.31 (2.54)
MD Velocity	1.94 (3.78)	2.37 (2.50)	-.41 (1.43)	0.01 (0.94)
MD Acceleration	2.44 (3.78)	2.31 (2.540)	-.04 (1.14)	-0.01 (0.85)
IOP	18.01 (5.30)	19.25 (5.14)	17.47 (3.82)	17.62 (3.10)
IOP Velocity	2.50 (3.87)	2.48 (3.02)	-.03 (2.00)	-0.10 (1.64)
IOP Acceleration	2.42 (3.56)	2.42 (2.74)	.05 (1.67)	0.01 (1.49)
PSD	12.57 (6.31)	4.93 (4.03)	8.20 (3.45)	3.37 (2.17)
Baseline MD	-5.35 (3.71)	-3.15 (2.60)	-5.35 (3.71)	-3.15 (2.60)
Baseline IOP	28.070 (5.70)	27.69 (5.04)	28.07 (5.70)	27.69 (5.04)
MD Change	-3.41 (6.56)	4.39 (3.52)	-3.21 (4.00)	0.84 (1.85)
IOP Change	-10.06 (7.53)	-8.44 (7.10)	-10.60 (5.21)	-10.07 (4.49)

Table 2.1: Summary statistics for Kalman filter and raw observations, separated by progressing and nonprogressing instances.

The final logistic regression models (an average of the 10-fold cross validation) are summarized in Table 2.2. Both models use the same set of final variables, however the odds ratios are larger for the model trained on Kalman filter estimates than the odds ratios for the model trained on raw observations. For instance, the odds ratio of PSD is 1.344 for the logistic regression model based on Kalman filter estimates and 1.107 for the logistic regression model based on raw observations. Analysis of the logistic regression fitted values, i.e. the estimated probability of progression, is presented in Table 2.3. It is important to note that the average estimated probability of progression of the Kalman filter progressing instances (0.738) is much higher than the average for the raw observations progressing instances (0.498). Additionally, the

difference between average fitted values of progressing and nonprogressing instances is much greater for the Kalman filter estimates (95% CI 0.611,0.676) than for the raw observations data set (95% CI 0.290,0.337). This increased difference of average fitted values between progressing and nonprogressing instances supports our hypothesis that Kalman filtered estimates allows for improved distinction of significant disease progression.

	Kalman Filter Estimates			**Raw Observations**		
Variable	Coefficient	Odds Ratio	P Value	Coefficient	Odds Ratio	P Value
Intercept	-3.492	0.030	<0.001	-1.966	0.140	<0.001
MD	-0.554	0.575	<0.001	-0.178	0.837	<0.001
MD Velocity	0.193	1.213	0.06	0.071	1.074	0.04
Baseline MD	0.813	2.256	<0.001	-0.009	0.991	0.86
PSD	0.297	1.345	<0.001	0.102	1.108	0.005
IOP	0.076	1.079	0.03	-0.024	0.976	0.03

Table 2.2: Final covariates for the Kalman filter and raw observations logistic regression models.

Model	**Progressing**	**Nonprogressing**	**95% Difference CI**
Kalman Filter Estimates	0.738	0.095	(0.611, 0.676)
Raw Observations	0.498	0.185	(0.290, 0.337)

Table 2.3: Fitted probabilities for progressing and nonprogressing instances from the Kalman filter and raw observations logistic regression models.

Finally, the ROC curve, Figure 2.1, illustrates that by first filtering the data using the Kalman filter model, we can achieve higher sensitivity and specificity than a model based on the raw observations. For example, if we select 95% sensitivity as our goal, we can obtain 83% specificity using the Kalman filter model but only 39% specificity using the raw observations model. At 90% sensitivity, the Kalman filter achieves 88% specificity while the raw observations model achieves 66% specificity. The mean (and variance) of the estimated AUC for the Kalman filter and raw observations models are 0.961 (0.002) and 0.889 (0.013), respectively. Hence, using the probability function generated via Kalman filter state estimates and GEE for logistic regression, we are able to more accurately classify patients and instances as

experiencing or not experiencing significant OAG progression.

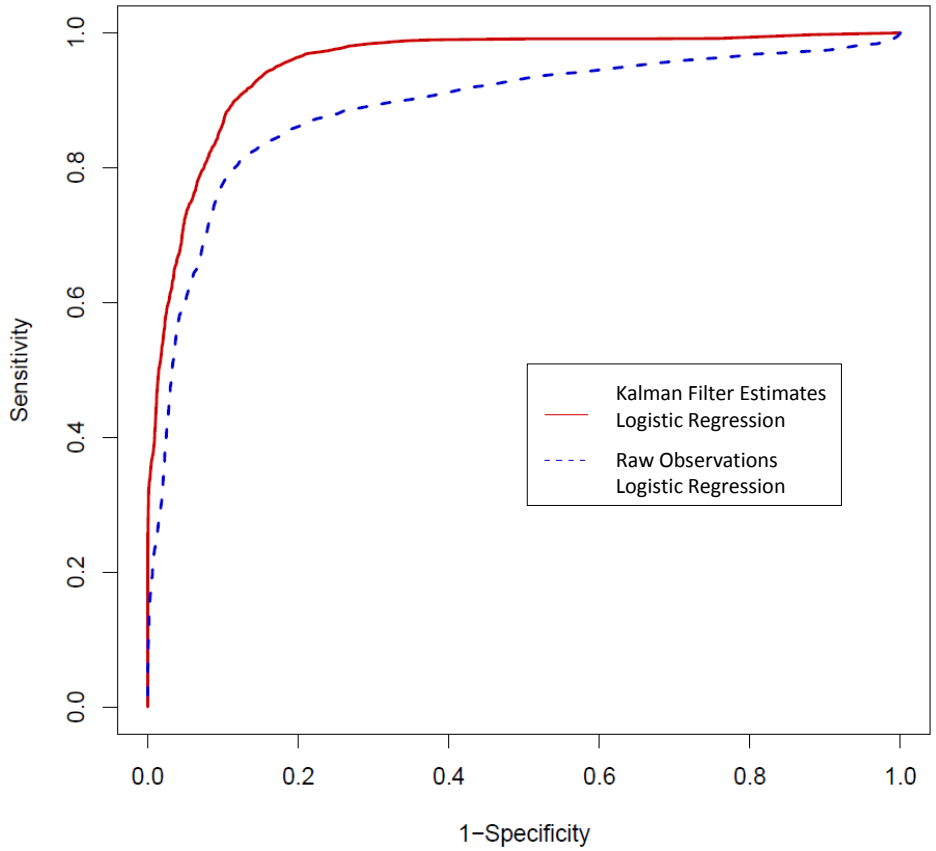

Figure 2.1: Receiver operating characteristic (ROC) curve for the Kalman filter and raw observations logistic regression models.

2.5 Discussion

Using Kalman filter forecasts in determinations of when patients with OAG should be observed by their physician required the development of a mapping from the filtered health metrics to a probability of progression. The application of GEE for logistic regression on Kalman filtered longitudinal observations of patients with OAG resulted in improved ability to identify significant glaucoma progression as compared to the model generated using the raw clinical trial data. The Kalman filter model is able to better detect relationships between health metrics and the more complex

disease progression definition than the logistic regression model using raw observations as inputs. We believe that as the progression definition becomes more heavily influenced by systematic process and measurement noise, a logistic regression model parameterized on Kalman filter estimates of the input will become increasingly more beneficial for detecting disease progression.

The methodology we present here takes advantage of state estimation and the linear system model of the Kalman filter, in conjunction with the marginal response of GEE, to improve the logistic regression model's ability to correctly classify patients. The Kalman filter model performs at higher specificity and sensitivities for significant disease progression classification due to the greater difference in mean fitted values (i.e. average estimated probability of progression) between progressing and nonprogressing instances. As we iterate through potential probability thresholds for classifying instances/patients as progressing or nonprogressing, the greater difference in mean fitted values creates a larger set of thresholds for which there are fewer false negatives and false positives. The lower rate of false negatives and false positives leads to improved sensitivity and specificity for the Kalman filter model at detecting significant glaucoma progression in comparison to the raw observations.

The difference of mean fitted values for the Kalman filter model is larger because of the greater in magnitude covariate coefficients and higher odds ratios of the model covariates. With higher odds ratios, each unit increase in a predictive covariate increases the probability of progression more greatly for the Kalman filter model than it does for the raw observations model.

The greater in magnitude covariate coefficients and higher odds ratios are explained by the linear system model the Kalman filter uses for state estimation. In the case of glaucoma, IOP decreases over time for treated patients and mean devi-

ation becomes more negative since VF loss cannot be reversed. The trend creates the larger difference in mean variable values as the number of measurements increases. The "noisy" nature of the raw observations creates fluctuation around this expected trend. Because the GEE approach is concerned with population-averaged (i.e. the mean response) variable, we expect covariate coefficients to be greater in magnitude when the difference between the mean variable value for progression and nonprogression instances increases.

The increased standard deviation of the mean value of the Kalman filter estimates is due to the Kalman filter's recognizing each patient's individual disease realization. As the Kalman filter updates the state estimates to reflect a patient's particular characteristics, that patient's trajectory becomes more dissimilar to the trajectories of other patients. The "noisy" raw observations mask these dissimilar trajectories which results in clustered mean variable values for progressing or nonprogressing instances.

Increased sensitivity and specificity of classification models improves clinical decision making by more accurately identifying significant disease progression. Clinicians who are able to correctly identify patients who experience significant glaucoma progression can make more informed decisions, such as improving monitoring schedules and improving treatment decisions. Additionally, the increased accuracy allows clinicians to utilize the statistical model without fearing high rates of misclassification.

Our proposed methodology is limited by the linear system dynamics model. Glaucoma progresses relatively slowly, thus changes in disease state can be estimated well by a linear model. For more rapidly progressing diseases, if the time between patient observations is sufficiently small, the disease progression mechanics can potentially be estimated by a linear dynamics model. The Kalman filter also assumes the state

estimates, noise and raw observations come from a Gaussian distribution. This assumption is reasonable within a range around the mean (2 standard deviations) for bounded variables, e.g. IOP.

The application of our methodology to CIGTS data is limited by the fact that this trial took place between 1993 and 2003. Since 2003, there have been many advances in the field of diagnostic testing for glaucoma including testing to check for damage to the retinal nerve fiber layer tissue using optical coherence tomography (OCT) and additional progression detection software on the visual field machines such as Guided Progression Analysis (GPA). In the future, we plan to use data from other sources to be able to integrate data from OCT and GPA into our models and progression definition.

2.6 Extension: Dynamic and Personalized Monitoring Schedules

Now that we have illustrated the use of Kalman filtering for noise reduction in statistical classification models, we next consider how the improved classification model can be used within the framework developed by [64] to better manage glaucoma patients. In particular, we considered how the classification model can inform monitoring decisions for individual patients. Monitoring decisions include deciding when a patient should be tested by the VF and IOP devices. When patients are tested too frequently, they undergo unnecessary testing and stress. However, when testing is too infrequent, clinicians may miss observing disease progression and miss the chance to intervene. The decision of how frequently to monitor a patient is directly affected by process and measurement noise. Therefore, our approach of combining noise reduction with statistical classification can enhance a clinician's ability to properly determine the monitoring frequency. In this section, our analysis has been expanded

by adding another source of randomized clinical trial data: the Advanced Glaucoma Intervention Study (AGIS). Together with CIGTS, these two clinical trials provide information on the disease dynamics for patients with early to advanced glaucoma. Since the AGIS data did not have information on light sensitivities of individual points of the eye, we were unable to continue using the HAP classification as part of our progression definition. Our revised progression definition is based on a validated drop in MD of at least 3 dB with respect to baseline MD [108].

2.6.1 Methods

The algorithm steps are illustrated in Figure 2.2. The algorithm uses information about the underlying population and the particular patients test results to estimate the patients true MD, IOP, PSD and the respective velocities and accelerations of these parameters. These key clinical values are then forecasted using the Kalman filter to compute confidence regions, centered around the estimate of the true values, at future time periods.

We then use the classification model derived in Section 2.3 to assess the probability, over the entire confidence region, that the patient will experience disease progression at the future time periods. If the maximum probability of progression for a confidence region exceeds a given threshold, the algorithm calls for a VF and IOP test at that future time period. Figure 2.5 provides a graphical representation of how the progression threshold determines the Time to Next Test (TNT). Once the VF and IOP tests are performed, the algorithm receives the new test measurements, the MD, PSD, IOP and the respective velocities and accelerations are updated using the Kalman filter. The algorithm then repeats this process of forecasting the patients disease parameters and computing the maximum probability of progression to determine when each patient should next be monitored.

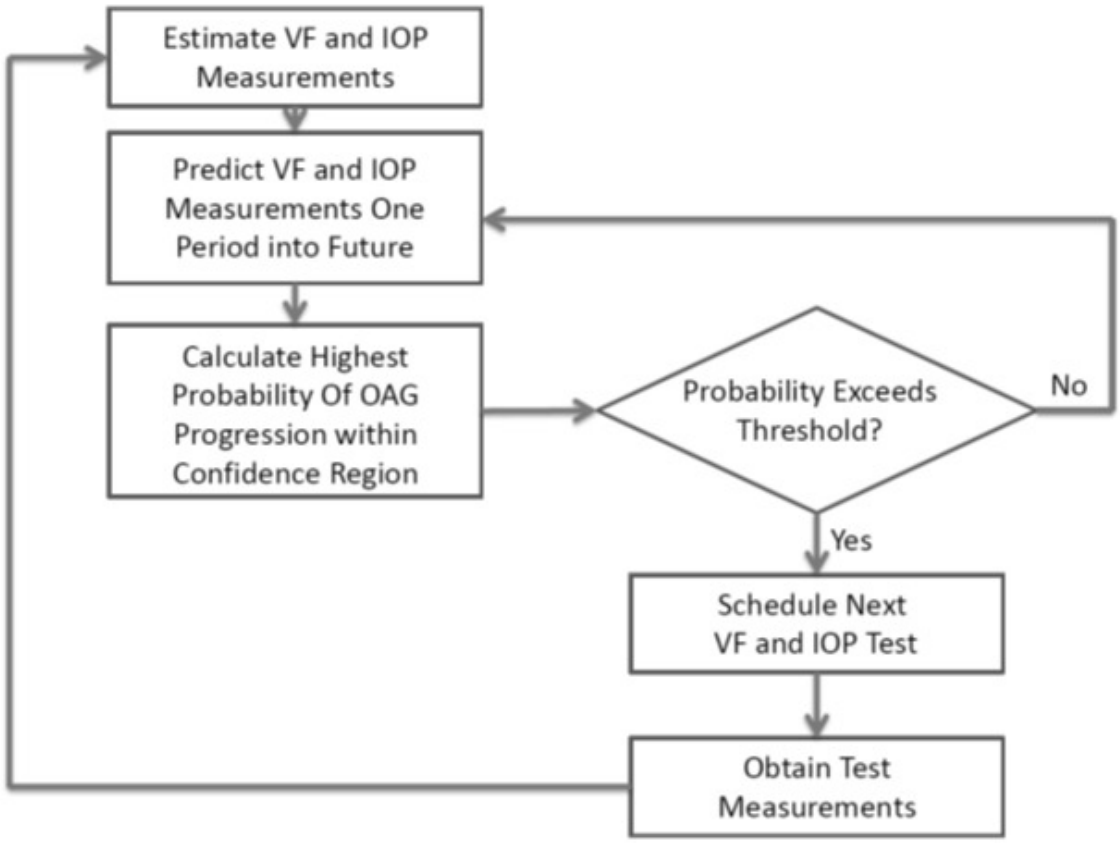

Figure 2.2: The Time to Next Test (TNT) algorithm flow diagram.

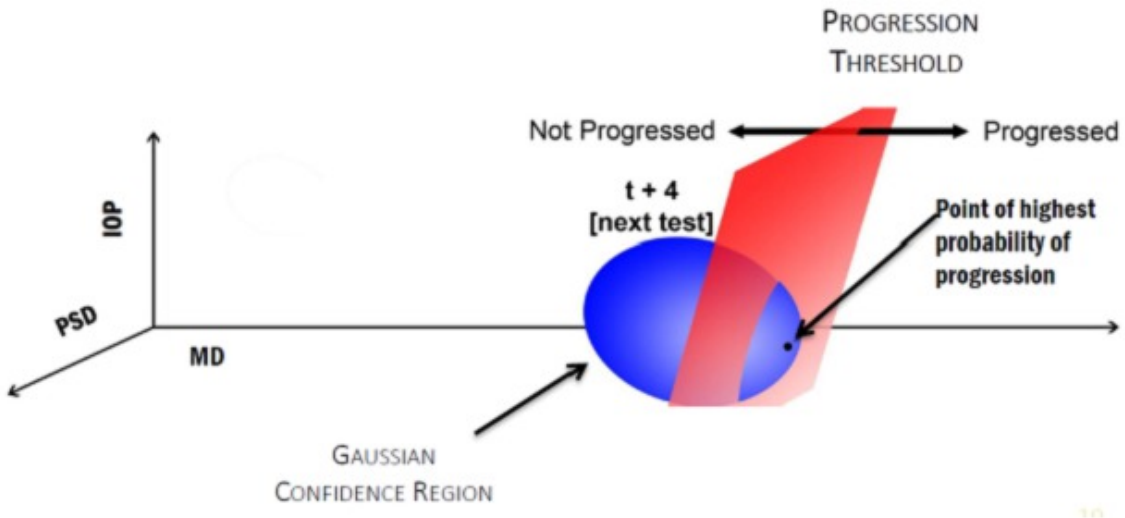

Figure 2.3: Probability threshold and the TNT algorithm

We compared the performance of our scheduling algorithm against 1 year, 1.5 year and 2 year fixed interval testing schedules for performing VFs and IOPs. To assess how well the algorithm performed relative to the fixed testing intervals, we compared: (i) the average number of examinations (VFs and IOP measurements) performed per patient per year; (ii) efficiency in testing (percent of instances where OAG progression was noted at the time a VF test and IOP measurement was scheduled); and (iii) diagnostic delay (average number of months that a patient's glaucoma progression went undetected between examinations). We used asymptotic values for efficiency (e.g. 50% for 1 year fixed) and diagnostic delay (e.g. 3 months for 1 year fixed) as the performance measures for fixed interval schedules. This algorithm was applied to the patient until a visit was scheduled on or after the date the patient first experienced glaucoma progression.

To validate and test our methodology, we divided the CIGTS and AGIS trial data equally into a training set (for parameterizing models) and testing set (for validating and testing the models). We randomly assigned CIGTS/AGIS participants to these sets to assure equal representation of both groups in the training and testing sets. We performed this randomization process 25 times and calibrated the Kalman filter for each randomization. The prediction error of the Kalman filter was consistently unbiased across the randomizations. We present here the numerical results of one of these randomizations.

2.6.2 Results

A total of 571 participants (571 eyes) with OAG met the study inclusion criteria. Table 2.4 presents a summary of the participants. Of these, 266 (47%) came from CIGTS and 305 (53%) came from AGIS. The mean (standard deviation) age of the study participants at baseline was 63.2 (10.9) years. The participants included 272

males (48%) and 299 females (52%). There were 263 whites (46%), 288 blacks (50%) and 20 were classified as some other race. Participants were followed in the trials for an average of 6.3 (2.8) years. The training dataset included 286 eyes of 286 patients and the testing dataset included 285 eyes of 285 patients. There was no statistically significant difference in the demographic characteristics, number of visits, or clinical parameters (mean MD, PSD, IOP) between individuals in the training and testing datasets ($p > 0.05$ for all comparisons) , except there were slightly more blacks in the training set than the testing set (154 vs. 134; $p = 0.05$).

		Training		**Testing**		**P Value**
Variable		n	%	n	%	
Number of Eyes		286		285		0.48
Number of Participants		286		285		0.48
Number from CIGTS		131	46	135	47	0.64
Number from AGIS		155	54	150	53	0.34
Sex	Male	135	47	137	48	0.57
	Female	151	53	148	52	0.40
Race	White	123	43	140	49	0.93
	Black	154	54	134	47	0.05
	Other	9	3	11	4	0.74
Total Number of Visits		3158		3227		0.89
Number of Progression Instances		163		166		0.59
Mean (SD) Number of Visits per Patient		11.0 (5.0)		11.3 (5.3)		0.4
Mean (SD) Age		64.2(10.9)		64.3 (11.0)		0.8
Kalman Filter Variables	Initial MD	-7.55 (3.74)		-7.65 (3.73)		0.91
	Initial PSD	6.49 (3.39)		6.41 (3.83)		0.54
	Initial IOP	17.61 (0.22)		17.7 (0.18)		0.28
	MD	-8.30 (2.04)		-8.27 (1.95)		0.96
	MD Velocity	-0.04 (0.21)		-0.03 (0.20)		0.23
	MD Acceleration	-0.01 (0.28)		-0.02 (0.26)		0.65
	PSD	6.58 (1.18)		6.70 (1.13)		0.69
	PSD Velocity	0.01 (0.13)		0.01 (0.12)		0.74
	PSD Acceleration	0.00 (0.19)		0.01 (0.17)		0.50
	IOP	17.43 (2.64)		17.14 (2.63)		0.23
	IOP Velocity	0.00 (0.29)		0 (0.31)		0.67
	IOP Acceleration	0.00 (0.42)		0.01 (0.41)		0.62

Table 2.4: Summary statistics for AGIS/CIGTS data used in TNT algorithm.

Figure 2.4 compares the average efficiency and diagnostic delay of the TNT algorithm and 1, 1.5, and 2 year fixed testing intervals. For the same average number of

tests as the 1, 1.5, and 2 year fixed testing intervals, the TNT algorithm achieved higher efficiency ($p<0.0001$ for all comparisons) and reduced diagnostic delay ($p=0.02$, <0.0001, and <0.0001 respectively) for detecting OAG progression. For example, when comparing the 1 year fixed testing interval against the TNT algorithm, for the same average number of tests (4.7 tests), the TNT algorithm increased efficiency by 29% and reduced diagnostic delay at OAG progression detection by 1.7 months. Table 2.5 compares the fixed yearly schedule against the TNT algorithm.

Performance Measures	**Fixed Interval Testing**	**TNT Algorithm**
Average Number of Tests per Year	1.00	1.12
Average Efficiency (%)	50	79
Average Diagnostic Delay (months)	3.00	1.29

Table 2.5: Performance measures for TNT algorithm and fixed yearly testing.

We also analyzed the subset of participants enrolled in both trials that experienced OAG progression as compared with those who never experienced glaucoma progression. Overall, 116 trial participants in the testing dataset were noted to have OAG progression and 169 did not exhibit progression. Among those in the testing dataset who progressed, the mean (SD) time from study enrollment to first record of OAG progression was 45.7 (23.4) months. Since efficiency and diagnostic delay assess the algorithm's ability to schedule follow-up tests at times when there was evidence of actual OAG progression, these performance measures were not applicable for the subset of participants who did not exhibit disease progression. The algorithm scheduled more tests per year for patients who were exhibiting OAG progression (1.3 tests per year) than others who were stable (1.0 test per year) ($p<0.0001$).

We further studied how the TNT algorithm performed on CIGTS and AGIS participants in the testing dataset separately. As one might expect, the TNT algorithm scheduled more tests for AGIS participants than CIGTS participants (1.3 average tests per year vs. 0.9 average tests per year; $p<0.0001$). The TNT algo-

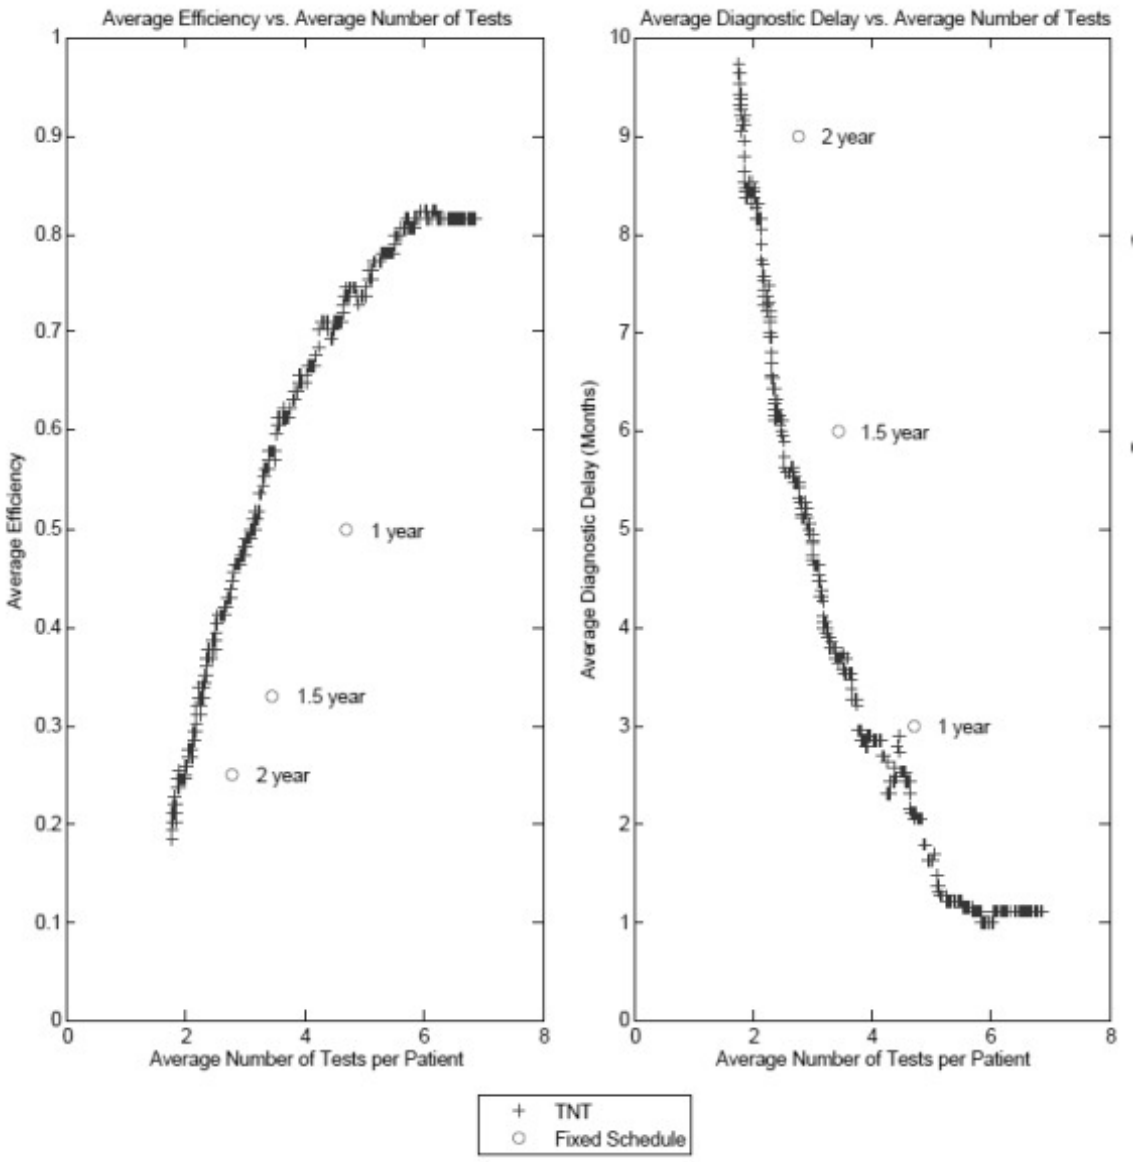

Figure 2.4: Comparison of the TNT algorithm and fixed-schedule performance measures.

rithm achieved marginally improved efficiency (83% vs. 71%; p= 0.06) for AGIS compared with CIGTS participants, and the efficiency at OAG progression detection for both groups were better than the efficiency achieved using 1 year fixed testing intervals (50%). Diagnostic delay at detecting OAG progression (1.0 months vs. 1.9 months; p= 0.09) was slightly shorter for AGIS participants, though this did not reach statistical significance.

Additional analyses were performed to see how well the TNT algorithm performed on black vs. white patients from the trials. We found that the TNT algorithm performed more tests on average for black patients than white patients (5.31 vs. 4.24; p=0.03). The TNT algorithm performed equally well in terms of efficiency and diagnostic delay (p=0.10 and 0.20, respectively) for black and white patients. Further analysis of the Kalman filter and the new logistic regression model can be found in Section 2.8.

2.6.3 Discussion

Comparing the output generated from the algorithm with fixed testing intervals of 1, 1.5, and 2 years, we show that the algorithm is capable of detecting OAG progression more efficiently and with reduced diagnostic delay compared with fixed interval schedules, without the need for additional tests. The model appears to work well for those with mild to moderate OAG (participants in CIGTS) as well as for those with more advanced disease (participants in AGIS), and performs well on the subset of trial participants who did and did not exhibit OAG progression, and forecasts well for white and black trial participants.

There are several advantages to using this approach to aid in the evaluating and monitoring of patients with OAG, rather than simply testing all patients at fixed intervals or relying on one's gestalt of how often to monitor a given patient. By incorporating data from a population of patients with OAG, the Kalman filter is able to identify and filter out systematic noise (e.g. measurement error, variability in test performance) that is known to exist in IOP readings and VF test results. Second, the Kalman filter makes use of data from sequential visits to account for the disease dynamics of each individual patient, and continually updates the model with new test results after each visit to determine the timing of future testing. Third, the

algorithm is scalable and can include additional data from structural tests such as optical coherence tomography or confocal scanning laser ophthalmoscopy, as well as other quantifiable data elements. Fourth, since there is presently no consensus on the optimal approach to define OAG progression, the model is flexible enough to be able to make predictions of progression using different definitions. Finally, the algorithm can be tailored by the eye care provider to be more or less aggressive in testing for disease progression. For example, the algorithm can be modified so that a clinician can choose to increase the threshold for detecting OAG progression for an 85 year old patient with early OAG who has multiple medical comorbidities, if she thinks this patient is unlikely to go blind from the disease, so as to not overburden such a patient with frequent tests. Alternatively, for a 40 year old monocular patient with severe OAG, the clinician might opt to lower the threshold so that the algorithm can identify the first hint of possible disease progression. From a societal perspective, the use of Kalman filter forecasting can improve the quality of care offered to patients by aiding in more timely identification of those who are exhibiting OAG progression and require additional treatment while simultaneously limiting patient burden and added costs of performing unnecessary testing.

There are several study limitations that need to be acknowledged. First, the types of parameters we were able to incorporate into the Kalman filter we developed were limited to those that were measured in the CIGTS and AGIS studies. Information that we would have liked to include in the algorithm but was not available from those trials includes pachymetry readings, optical coherence tomography measurements, and other quantifiable measures of the optic nerve or retinal nerve fiber layer. In the future we hope to obtain access to datasets that longitudinally capture information on these parameters so we can refine our algorithm, which should enhance its

ability to identify persons who are at increased risk of OAG progression. Second, we have yet to test this algorithm on other groups of patients such as those with ocular hypertension, those with early pre-perimetric glaucoma, those with other forms of glaucoma, and those who underwent incisional glaucoma surgery. Further validation is necessary to determine how well the algorithm predicts disease progression and need for monitoring in these groups. Third, the timing of the follow-up examinations in AGIS and CIGTS restricted our algorithms scheduling decisions to no more frequently than every 6 months. If follow-up examination data for smaller time windows, e.g. every 1 month, were available, our algorithm could make scheduling decisions as often as every month. As we shorten the time interval allowed in scheduling (e.g. 6 months to 1 month), we expect the algorithm to achieve higher efficiency and lower diagnostic delay. In particular, this would have large gains in the improvement of our TNT algorithm for diagnostic delay. The exact gains cannot be known until we have tested our TNT algorithm on data collected at the higher frequency of every 1 or 3 months. And lastly, patient adherence to prescribed medications is likely higher for participants in AGIS and CIGTS compared to those routinely care for in clinical practice. When applied to patients seen in clinical practice, the increased IOP variability due to lower medication adherence would likely decrease the predictive capability of the Kalman filter. When we further validate the model on another sample of patients who were not enrolled in a clinical trial we will be able to explore this further.

2.7 Conclusion

In this research, we applied a linear system dynamics model approach, using a Kalman filter, to estimate true measurement values for variables which have both

measurement and process noise. Filtering techniques are important for true measurement estimation for medical decision making and have been shown to result in improved significant disease progression classification when utilizing GEE for logistic regression with repeated measures data, as demonstrated in our modeling of OAG progression dynamics. Due to process and measurement noise, only after having seen future observations can clinicians retrospectively assess whether true progression has occurred. Logistic regression models that directly consider those noises allow for the prospective calculation of the probability of experiencing progression. Furthermore, for complex progression definitions, logistic regression enables a reduction in the number of variables to consider, which is important in guiding clinical decisions. We showed that our statistical classification approach reduces misclassification error via noise extraction. Next, we extended this research to study how an improved classification model can be used to better monitor glaucoma patients. We combined the classification model with the Kalman filter forecasts to develop dynamic and personalized monitoring schedules that outperform the current practice of fixed interval schedules. While this methodology has been applied to OAG, the techniques are applicable to many other chronic diseases, particularly those diseases whose dynamics can be modeled effectively by a linear system and whose biomarkers can be reasonably approximated by a Gaussian distribution.

2.8 Appendix

Table 2.6 presents the coefficients, standard errors, and p-values of the covariates incorporated into the logistic regression, which we then used to assess the probability of OAG progression for each patient. As expected, patients with more advanced glaucoma as captured on perimetry (a more negative MD or a more positive PSD) had a higher probability of progression compared to those with less advanced disease. In the regression model, each of the covariates in Table 2.6 was found to be significantly associated with OAG progression ($p < 0.04$ for each covariate).

Covariate	**Coefficient**	**Standard Error**	**P Value**
Intercept	-6.004	0.723	<0.001
Mean Deviation (dB)	-0.057	0.017	0.001
Mean Deviation Velocity (dB per month)	-4.054	0.666	<0.001
Mean Deviation Acceleration (dB per month^2)	-1.183	0.326	<0.001
Baseline Pattern Standard Deviation (dB)	-0.162	0.078	0.039
Pattern Standard Deviation (dB)	0.154	0.075	0.039
Age (years)	0.026	0.103	0.013

Table 2.6: Final covariates for logistic regression in TNT algorithm.

To validate the fit and predictive ability of the Kalman filter for assessing OAG progression, we calculated the 95% confidence intervals for the mean prediction errors of MD, PSD, and IOP and their respective velocities and accelerations across all study participants in the testing dataset. Errors were calculated at various prediction lengths (6 months, 2 years, and 5 years into the future). Table 2.7 shows that the mean differences between the Kalman filter predictions and the observed values from the trials were close to zero across various prediction lengths ($\alpha = 0.05$), supporting the accuracy of the Kalman filter predictions.

We also compared the observed values of MD from each clinical trial participant in the testing dataset against the filtered and predicted values of MD generated by the Kalman filter. The Kalman filter forecasts one period ahead and updates the

	95% Confidence Interval		
Variable	6 Months	2 Years	5 Years
Present MD	(-0.1493, 0.0329)	(-0.1294, 0.1361)	(-0.2979, 0.3882)
MD Velocity	(-0.0163, -0.0003)	(-0.0175, 0.0024)	(-0.0361, 0.0000)
Present PSD	(0.0163, 0.1227)	(0.1308, 0.2880)	(0.4675, 0.8394)
PSD Velocity	(0.0065, 0.0168)	(0.0039, 0.0166)	(-0.0028, 0.0195)
Present IOP	(-0.2487, 0.0121)	(-0.2751, 0.0792)	(-0.6454, 0.0519)
IOP Velocity	(-0.0033, 0.0194)	(-0.0117, 0.0185)	(-0.0365, 0.0175)

Table 2.7: Prediction error for the Kalman filter.

forecasts with the clinical observation for that period to obtain the filtered estimate of MD at each sequential trial visit. Predicted MD values are those obtained from the Kalman filter without incorporating future clinical observations. To illustrate the Kalman filters forecasting ability, we present in Figure **??** four study participants, two of whom exhibited OAG progression and two which experienced no progression during their enrollment in one of the clinical trials. We also estimated 90% confidence intervals for the predicted values toward the end of each participant's enrollment in the clinical trial. We chose the narrower 90% confidence intervals for the predicted values to demonstrate how strong the predictive power of the Kalman filter actually is. Since all observations fell well within the 90% confidence intervals, the observations would also fall within the wider 95% confidence intervals. We found that at all future time points, the Kalman filter forecasts for MD were close to the observed MD values obtained when the participant took the test during the clinical trial; our confidence intervals for predicted MD fully encompassed the observed MD values, even 3.5 years into the future. Similar analyses were performed on all patients in the testing set for PSD and IOP. Figure 2.6 shows an example of how the algorithm forecasts future PSD and IOP measurements.

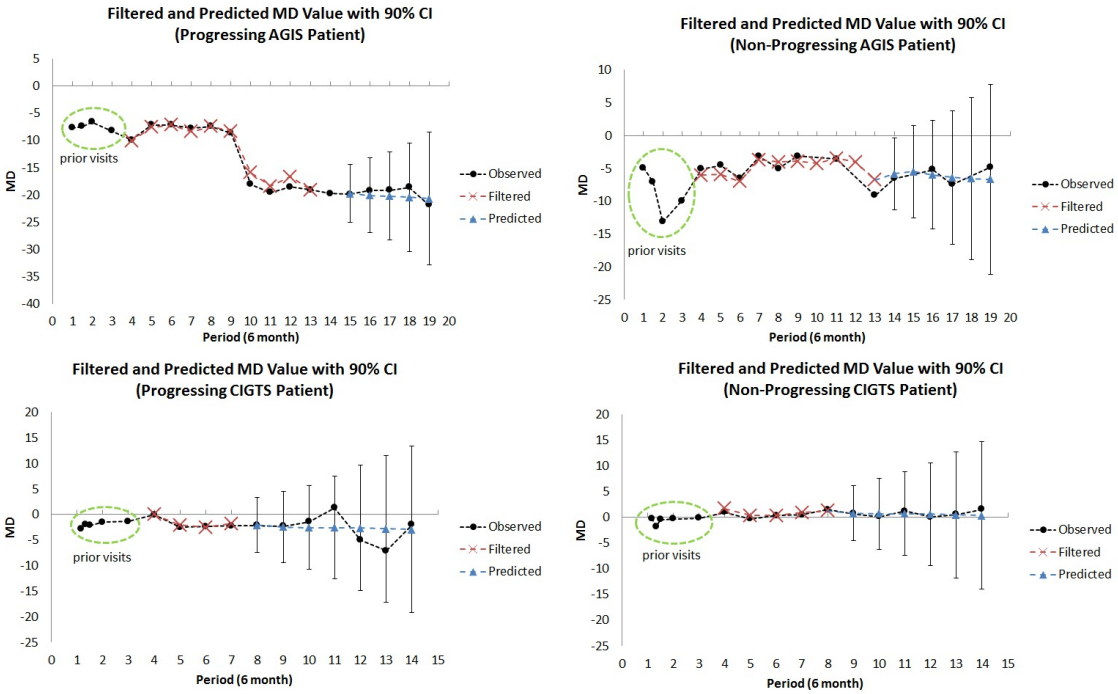

Figure 2.5: Kalman filter trajectories of mean deviation (MD).

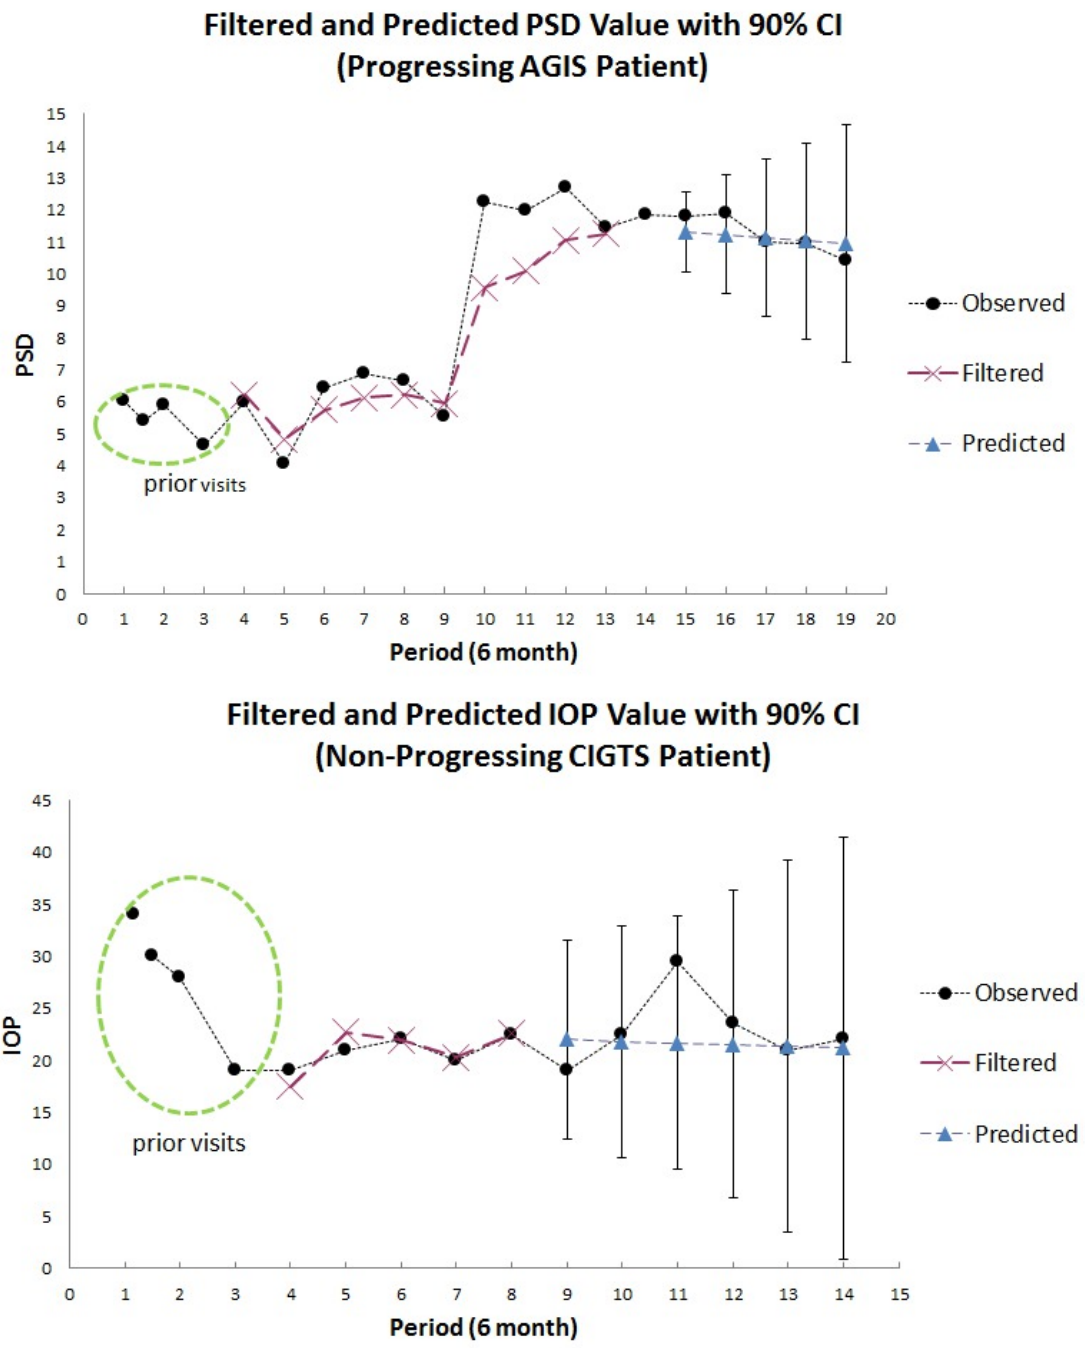

Figure 2.6: Kalman filter trajectories of pattern standard deviation (PSD) and intraocular pressure (IOP).

CHAPTER III

Treatment Planning for Cardiovascular Disease

In this chapter we discuss dynamic programming formulations for the optimal long-term treatment of patients at risk for cardiovascular disease events, e.g. heart attack and stroke. This research focuses on analyzing patient heterogeneity, which includes disease risk factors and treatment consumption, in order to personalize the treatment policies. Our work establishes the modeling framework for incorporating patient heterogeneity, via conditional value-at-risk, to tailor the reward structure of the dynamic program to the particular characteristics of the patient. The properties of the optimal and personalized treatment policies are studied, and the policies are compared against current treatment guidelines using observational patient data. Next, we extend our formulation to consider a different objective function and use the new model to answer key clinical questions: (1) how does misestimation of treatment disutility impact treatment decisions and patients and (2) how do time-varying treatment benefit and disutility influence the timing of initial treatment.

3.1 Background

Heterogeneity in the stochastic outcome of control actions translates to variability in rewards/costs accrued in a sequential decision making framework. Within a medical decision making context, this heterogeneity occurs when differences in patient

characteristics result in different outcomes for the same treatment. Heterogeneity may arise due to variability in patient characteristics, how biological factors affect disease response to treatment, as well as differing patient perception, risk sensitivity, and other factors that impact the medication consumption of the patient. Failure to identify and incorporate sources of heterogeneity may lead to improper specification of the optimal control policy and suboptimal performance when used in practice. This suboptimal performance can have strong and immediate consequences in medical decision making. To address heterogeneity, we propose a general method which utilizes a state-space inclusion of heterogeneity parameters to compute the expectations of truncated distributions. This robust approach enables both the identification and exploitation of heterogeneity in the development of optimal treatment policies and the understanding of how changes in heterogeneity parameters influence health outcomes.

Heterogeneity in control action outcomes and the associated consequences are especially important in the context of lifetime treatment decisions for patients with chronic disease. In the clinical setting, the observed variability in disease response to cancer therapies [28] and HIV treatments [112] have become popular examples of the need for tailoring treatment decisions to a patient's particular characteristics [48]. Personalized medicine poses a challenge to clinicians [156] and an opportunity for operations researchers to exploit these medical findings of heterogeneity to improve the performance of treatment policies. When designing personalized treatment policies, it is also important to consider the side-effects associated with the intensity of the prescribed medications. For this reason, the total treatment intensity (or resource utilization) should be constrained to prevent harm from increased side-effects under the optimal treatment policies.

We focus on modeling heterogeneity parameters and the corresponding patient-specific treatment outcomes (e.g. disease response distributions) in a stochastic control problem for patients at risk of coronary heart disease (CHD) events. CHD is the leading cause of death for men and women worldwide, and it caused approximately 1 out of 6 deaths in the United States in 2009 [77]. CHD is caused by atherosclerotic plaques in the heart's arteries. These plaques can lead to narrowing of the arteries, but even more importantly, are prone to rupture. When plaques rupture, a blood clot can form in the artery and block blood flow to the heart muscle downstream. A heart attack (also called a myocardial infarction) occurs when heart muscle dies, causing irreversible damage [128]. One of the most important and controllable risk factor for cardiovascular disease is high blood pressure, i.e. hypertension [155]. More than 42 million Americans have hypertension that warrants treatment [38]. The U.S. guidelines for treating hypertension follow a treat-to-target (TTT) strategy designed by the Joint National Committee (JNC) on Prevention, Detection, Evaluation, and Treatment of High Blood Pressure [32]. Clinicians are advised to initialize non-diabetic patients on a single antihypertensive medication if their systolic blood pressure (SBP) is between 140 mmHg and 159 mmHg (stage 1 hypertension). If their SBP is higher (stage 2 hypertension), clinicians are advised to start patients on a two drug combination. Guidelines recommend changes in dosages of drugs or the addition of drugs until the goal blood pressure of 140 mmHg (130 mmHg for diabetic patients) is achieved. Aside from diabetes status, the TTT guidelines are not sensitive to the patients particular disease risk factors which determine the likelihood that the patient will experience CHD events. Furthermore, only 53.1% of persons with hypertension have achieved the target blood pressure [59]. Patient adherence to prescribed hypertension treatments is, at least in part, to blame for poor

hypertension control since controlled blood pressure is highly correlated with treatment consumption (3.44 odds ratio of good blood pressure control among adherent patients compared against nonadherent patients) [83].

The poor blood pressure control and high incidence of CHD death under the TTT strategy indicates a need for evaluating changes to clinical management policy. Improving health outcomes for hypertensive patients at risk for CHD events requires a joint effort on the part of treatment guideline policymakers and health providers. Treatment guideline policymakers should tailor guidelines to reflect patient heterogeneity, while health providers may consider interventions aimed at improving patient adherence to the guidelines. To assist the efforts to improve the management of hypertensive patients, we develop a model which incorporates heterogeneity in CHD risk factors and treatment consumption. Figure 3.1 provides an overview of the key features and decision makers involved in improving CHD patient health outcomes. Within this context, we study three potential changes: (1) changes in treatment guidelines under the current resource utilization, (2) intensification of resource utilization, and (3) interventions aimed at improving treatment consumption.

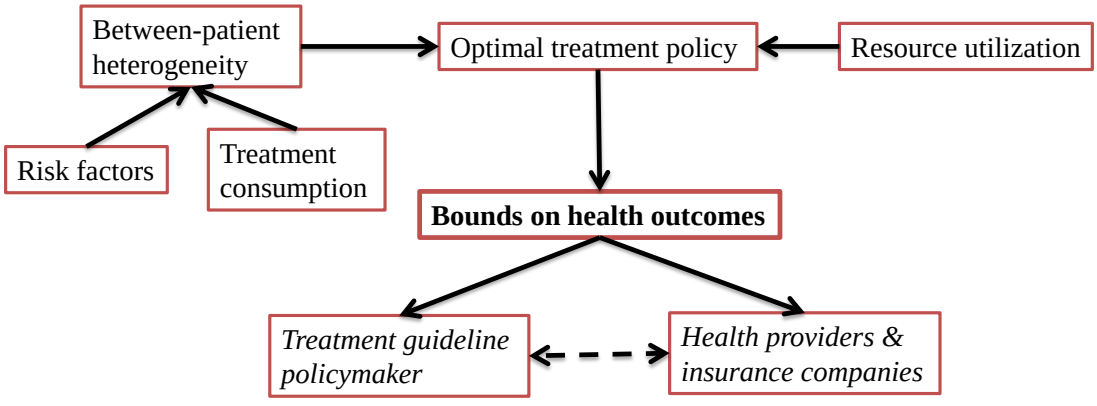

Figure 3.1: Decision makers involved in improving CHD patient health outcomes.

We start by formulating the lifetime hypertension treatment problem as a single patient, finite-horizon, discrete-time Markov decision process (MDP). We apply this model to each patient within a representative sample of the population. We characterize heterogeneity between patients into two sources: risk factors (such as age, sex, and diabetes) that directly impact disease progression, and factors (such as patient perception, risk sensitivity, education, and socioeconomic status) that impact treatment consumption. The model accounts for heterogeneity in risk factors and treatment consumption, as well as resource utilization, when deriving the optimal policy that minimizes the patients expected weighted number of CHD events over the planning horizon. To understand the role of treatment consumption, we consider adherence to medication as a motivating example of treatment consumption in the context of hypertension treatment planning. The MDP uses conditional value at risk (CVaR) to model patient-specific treatment outcomes which are affected by heterogeneity. CVaR is a coherent risk measure that isolates subsections of the treatment outcome distribution in order to personalize the expected benefit from treatment.

We then consider the perspective of a hypertension treatment guideline policymaker, such as the Joint National Committee that is tasked with studying hypertension and making treatment recommendations. This committee may be comprised of physicians, statisticians, epidemiologists, hospital administrators, and government agency workers. This treatment guideline policymaker wants to determine optimal and personalized treatment policies that account for heterogeneity across patients while ensuring a fair comparison between the TTT strategy and new optimal treatment policies. The guidelines may directly tailor treatments to each patients estimated adherence or, alternatively, the guidelines may assume perfect adherence for the entire population. The evaluation of the marginal benefit of resources is also of

interest to the treatment guideline policymaker.

Finally, we investigate clinical policy changes from the perspective of a health provider or insurance company that is considering implementing treatment consumption improving interventions (e.g. pharmacist intervention [134], cultural education [18], and messaging systems [68]). The provider wants to establish an upper bound on the marginal benefit of improving the patient's adherence to medication.

By using a unified framework for decision making and cardiovascular disease dynamics, treatment guideline policymakers and health providers can coordinate their efforts to maximize patient health. While new, optimal treatment guidelines are implemented, health providers and insurance companies can pursue adherence improving interventions given the knowledge of the marginal benefit of improving adherence. Since the U.S. Department of Veterans Affairs (VA) serves as both health provider and insurance company to the men and women of the U.S. armed services, we parameterize our model using longitudinal data from a nationwide, multi-center VA data set.

We consider the following to be our main contributions: (i) we develop a mathematical model that incorporates risk factor and treatment consumption heterogeneity to determine optimal hypertension treatment guidelines for patients at risk of CHD; (ii) we directly model resource utilization to fairly compare treatment policies and to compute the marginal benefit of resources; (iii) we utilize our generalizable modeling framework to compute the marginal benefit of improving treatment consumption on patient health outcomes; (iv) we derive analytical results relating heterogeneity parameters and resource utilization with the optimal value function and the optimal treatment policy; (v) we provide numerical results from the application of our model to nationwide longitudinal data which support our analytical findings; and (vi) we

compare U.S. guidelines to the optimal hypertension treatment policy given by the MDP formulation and show the health improvement achievable by the MDP policy.

3.2 Literature Review

3.2.1 Treatment Models

Systematic reviews of Markov decisions processes and their applications to treatment decisions can be found in [3] and [136]. While many MDPs for treatment decisions focus on small action spaces, such as binary decisions (e.g. wait to initiate therapy or initiate therapy) including the work by [2], [4], [5], [31], [133], [81], [13], [139], and [141], our model utilizes a larger action space to capture the intensity of treatments. Previous work with an expanded action space include: the investigation of multiple therapies for HIV control [140]; treatment and testing decisions for hepatitis C [49], [76]; optimal dosage level of a given treatment over a series of treatment sessions [58]; and dosage decisions for anesthesia infusion decisions [70]. Our work differs from this research in that we explicitly model and update estimates of patient heterogeneity, and we model directly a patient's resource availability to study the impact of resource limitations on the optimal treatment policy.

Researchers have also investigated treatment decisions for patients at risk of CHD. [84] and [45] developed a MDP formulation of statin therapy initiation for patients with type 2 diabetes who are at risk of CHD that maximizes quality-adjusted life years (QALYs) with considerations of costs. [63] used an infinite-horizon POMDP approach to decide medical therapies for patients with ischemic heart disease. The objective was to minimize the cumulative cost of treatment, represented as a combination of economic, quality of life, and health costs. [102] used a MDP formulation to determine blood pressure and cholesterol treatment decisions for patients at risk of CHD. The action space is comprised of wait/initiate decisions for a set of blood pres-

sure and cholesterol treatments. She used a bi-criteria objective function of QALYs and costs using a willingness-to-pay factor.

While we consider the research of [102] and [58] to be most closely related to our work, our inclusion and analysis of heterogeneity parameters in our sequential decision making framework differentiates our research from all of the previously mentioned work. In addition to the treatment consumption heterogeneity parameter, our work differs in that our action space is a large collection of treatment intensities, and we allow for the discontinuation of treatments. Furthermore, the explicit modeling of resource utilization also distinguishes our work from the previous research in treatment modeling. Through direct inclusion of resource availability in the model, we are able to perform marginal-benefit analysis on the resources.

3.2.2 Conditional Value at Risk

One key aspect of our work is the utilization of conditional value at risk (CVaR) for computing treatment effects and capturing the impact of patient heterogeneity in treatment consumption. CVaR is a risk measure popularized in the financial literature which calculates the expected value of a random variable over a restricted region of the random variable's domain [126], [127]. Let X by a random variable with cumulative distribution function $F(z) = P(X \leq z)$, where X has the meaning of a gain (e.g. improvement in health due to treatment). Define the value at risk ($VaR_\alpha(X)$) as the lower α-percentile of the random variable X. For $\alpha \in (0, 1]$, VaR is computed as:

$$VaR_\alpha(X) = \min\{z \mid F(z) \geq \alpha\} \tag{3.1}$$

CVaR of X with level α, denoted $CVaR_\alpha(X)$, is then the mean of the generalized

α-tail distribution [135], i.e.

$$CVaR_\alpha(X) = \int_{-\infty}^{\infty} z dF^\alpha(z) \tag{3.2}$$

where

$$F^\alpha(z) = \frac{\alpha - F(z)}{\alpha} \mathbb{1}\{z \leq VaR_\alpha(X)\} \tag{3.3}$$

For continuous distribution functions, we can simply write

$$CVaR_\alpha(X) = \mathbb{E}[X \mid X \leq VaR_\alpha(X)] \tag{3.4}$$

Popularly, the parameter α measures risk-sensitivity. Risk-sensitivity can play an important role in medical decision making since some patients will be risk-averse (i.e. place greater weight on worst-case outcomes) and others may be risk-seeking (i.e. place greater weight on best-case outcomes) depending on the severity of their disease condition. In our application to hypertension treatments, we consider an equivalent interpretation of α, where α is our treatment consumption heterogeneity parameter for the patients adherence to therapy. In this sense, $\alpha = 1$ implies full adherence to therapy, while $\alpha = 0$ implies complete nonadherence to therapy.

Attention to the role of patient adherence to therapy in sequential treatment decision models is relatively new in the literature. [164] modeled prostate biopsy referral decisions, but assumed imperfect adherence was implicitly captured by the random nature of prostate-specific antigen levels. [141] considered a patient-dependent multiplier for adherence that leads to proportionally higher (lower) expected lifetime upon initiating therapy if the patient is more (less) adherent than the average patient. [103] modeled adherence as a state variable in their statin treatment model. However, adherence was limited to four categories and it was assumed either the patient had a fixed adherence state or evolved in accordance to a transition matrix.

The patient-dependent multiplier and the categorical adherence levels are special cases of our CVaR approach. And [12] highlighted that mammography screening has adherence issues, but left such extensions to future work.

Unlike previous research, our approach is not restricted by the assumptions on the effect of treatment consumption on treatment outcomes. Our generalizable approach allows for discrete or continuous and parametric or empirical distributions for treatment outcomes. Our method explicitly connects the treatment consumption heterogeneity parameter α for patient adherence to a general distribution $\mathcal{F}_\tau$ used to compute immediate costs in the MDP formulation. Through CVaR, we can effectively model any assumed treatment consumption effects. For example, we can let $\mathcal{F}_\tau$ be a uniform distribution, $U_\tau(0, M)$, to capture patient-dependent multipliers and proportionality, where $M < \infty$ is the maximum disease response from therapy τ. Letting $\alpha = 0.5$ be the average patient adherence, then $\alpha = 1$ would be a patient with maximal adherence, and $\alpha = 0$ would be a patient who never complies with therapy τ (and therefore receives no benefit from the treatment). Other parametric distributions may also be used; for instance, we may consider a Gaussian distribution when the rate of improvement in disease response (due to increased adherence) is large when starting from moderate adherence levels, and the rate is small when starting from high adherence levels.

We may also consider empirical distributions, both discrete and continuous, for the treatment outcome. Discrete empirical distributions arise when treatment consumption is discretized or categorized into intervals and disease response is estimated from data. For example, [117] discretized adherence to HIV therapy into intervals: $\geq$ 95%, 90-94.9%, 80-89.9%, 70-79.9%, and $<$ 70% and the percentage of patients with virologic failure was computed: 21.7%, 54.6%, 66.7%, 71.4% and 82.1%, respectively.

For this discrete distribution, we can define probability mass functions on x_α that lead to x_{95}=21.7; x_{90}=54.6; x_{80}=66.7; x_{70}=71.4; and x_0=82.1. Similar construction can be done for empirical continuous distributions, as in the case of modeling the effect of adherence and drug resistance [15].

The previous work in the field that studied adherence are all special cases of CVaR. Our approach is generalizable to any of the assumed relationships between treatment consumption and disease response in the literature (see [74] for a review of this medical literature), which makes optimal control policies based on CVaR computations more widely applicable.

CVaR in stochastic control has arised primarily in two cases: (1) as a candidate risk measure in risk-sensitive Markov decision processes [100], [116], [19], and (2) as a risk-based constraint in a Markov decision process [20], [82]. The first case seeks to minimize a risk measure u, such as CVaR, of total cost: $\sum_{n=0}^{N} E[u(c_n)]$, where c_n is the cost incurred when there are n decisions remaining. While the second case uses CVaR to prevent low-likelihood and undesirable outcomes. Our approach differs in that we are concerned with minimizing total cost (e.g. expected total number of CHD events) which is a function of CVaR random variables: $\sum_{n=0}^{N} E[c_n(u)]$. Rather than using CVaR as the overall objective, we are using CVaR implicitly in our objective function via computation of costs and the differentiation of patients.

3.3 MDP Model Formulation

We formulate the lifetime hypertension treatment problem as a single patient, discrete-time, finite-horizon Markov decision process. Our objective is to minimize the patient's expected weighted number of CHD events over the planning horizon while considering sources of patient heterogeneity and resource utilization. Formally,

the notation for the MDP is as follows:

- N: total number of decision epochs in the planning horizon
- $\tau \in A$: treatment τ from the set of all treatment options A.
- Θ: treatment consumption heterogeneity matrix ($|A| \times N$) comprised of the patient's treatment consumption heterogeneity parameters with elements $\alpha_{ij} \in [0,1], i = 1, \ldots, |A|, j = 1, \ldots, N$.
- $K \in \mathbb{R}^+$: patient's resource amount available over the planning horizon.
- $s \in S$: multidimensional state of the patient represented by the number of remaining decision epochs in the planning horizon, the patient's treatment consumption heterogeneity parameter, the remaining resource available, and risk factors including demographic information for the patient (e.g. age, sex, diabetes status, smoking status), his/her pretreatment SBP, HDL, TC, and the patients health state. Let $n = 0, \ldots, N$ be the number of remaining decision epochs (also a proxy variable for the patients age); $y \in \{0, \ldots, K\}$ denotes the remaining resource available for expenditure; p represents patient demographic information necessary for computing CHD risk, $b \in \mathbb{R}^+$ denotes the pretreatment SBP; $l \in \mathbb{R}^+$ denotes the HDL; $c \in \mathbb{R}^+$ denotes the TC; and $h \in \{1, 2, 3, 4\}$ denotes the patient health state (no CHD event; survived a CHD event; died from a CHD event; died from non-CHD cause, respectively). Thus, $s = (n, \Theta, y, p, b, l, c, h)$ is a particular state vector capturing risk factor and treatment consumption heterogeneity.
- $r(s) \in [0,1]$: patient's one-epoch pretreatment risk of a CHD event when in state s.

- $b_{min} \in \mathbb{R}^+$: minimum allowable SBP.
- $A(b_{min}; y) \subseteq A$: set of available treatments sorted by their treatment intensity (computed based on the dosage and number of drugs prescribed) and restricted by the minimum allowable SBP b_{min} and amount of resource remaining y.
- $\gamma_\tau \in \mathbb{R}^+$: resource cost (measured by the intensity of the treatment) associated with selecting treatment τ.
- $\mathcal{F}_{\tau,\int}$: cumulative distribution function for the SBP reduction of treatment τ when in state s.
- $x_{\alpha,\tau}(s) \in \mathbb{R}$: CVaR SBP reduction from treatment τ with treatment consumption heterogeneity parameter α when in state s.
- $D(s) \in [0, 1]$: relative-risk factor when in state s.
- $\phi(s) \in [0, 1]$: probability of death from non CHD cause given the patients post-treatment CHD risk when in state s.
- $\rho(s) \in [0, 1]$: probability of death from a CHD event when in state s, given that the patient had a CHD event (i.e. a subsequent CHD event), independent of the treatment selected.
- $\zeta \in [0, \infty)$: pretreatment CHD risk multiplier for subsequent CHD events.
- $r_{\alpha,\tau}(s) \in [0, 1]$: one-epoch risk of a CHD event after taking treatment τ when in state s with treatment consumption heterogeneity parameter α.
- $p_\tau(s, s') \in [0, 1]$: probability that the patient transitions from state s to post-state s' after selecting treatment τ.
- $J(s)$: immediate cost accrued when the patient is in state s.

- $V(s) \in \mathbb{R}^+$: optimal value function that gives the minimized expected weighted number of CHD events when in state s. By construction, $V(n = 0, \ldots) = 0$.

At the beginning of each decision period, we assume that a measurement of the patient's risk factors (e.g. demographic information, pretreatment SBP, high density lipoprotein (HDL), and total cholesterol (TC)) is obtained. Using a risk calculator, such as Framingham [39], we estimate the patient's one-period pretreatment risk of a CHD event. We consider a treatment consumption heterogeneity matrix, Θ, comprised of α elements for each decision period and treatment action combination. In the context of adherence to therapy, this matrix allows for age- and treatment-specific adherence levels. We consider the α parameter to be a summary variable for the patients treatment consumption, which likely depends on other variables such as socioeconomic status and education. Furthermore, while Θ is an element of the state vector s, we repeat for emphasis the dependency on a particular α (an element of Θ) for some variables in our notation. Figure 3.2 provides an overview of the steps taken within each decision epoch of the MDP model. At each decision epoch, the clinician prescribes a hypertension treatment for the patient to follow until the next decision epoch. Given the patient's treatment consumption α and the SBP reduction function, $\mathcal{F}_{\tau,f}$, for each possible treatment τ, we compute the CVaR of the SBP reduction for each treatment, $x_{\alpha,\tau}(s)$, to determine the true SBP reduction under the given level of adherence. Only those treatments whose CVaR reduction in SBP would not bring the patient's SBP below the minimum allowable SBP, b_{min}, are feasible treatments in a given decision epoch.

We compute the CVaR SBP reduction $x_{\alpha,\tau}$ for treatment τ in accordance to equation (3.2), i.e. $x_{\alpha,\tau} = \mathbb{E}[X \mid X \leq \sup \mathcal{F}^{-1}_{\tau,s}(\alpha)]$.

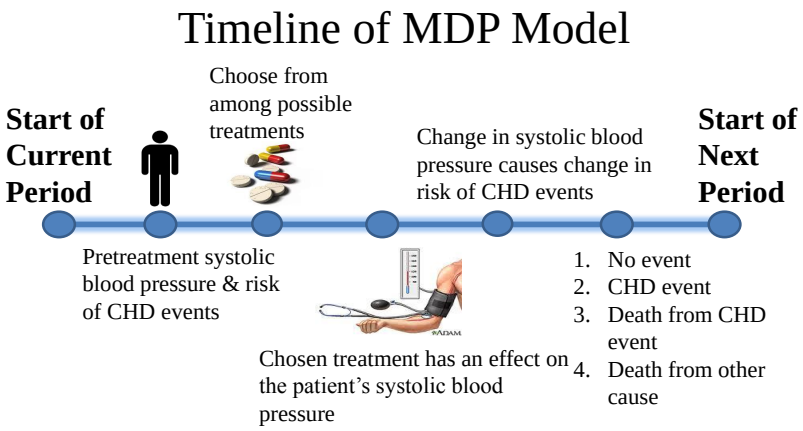

Figure 3.2: Hypertension treatment timeline.

While $\mathcal{F}_{\tau,s}$ is parameterized as a cumulative distribution function, this function is not used as a probability distribution of outcomes. Rather, $\mathcal{F}_{\tau,s}$ serves as a mapping for how changes in treatment consumption (e.g. $\alpha = 0.56$ to $\alpha = 0.60$) affect the SBP reduction from treatment, $x_{\alpha,\tau}$. We utilize the shape of the distribution function $\mathcal{F}_{\tau,s}$ and the CVaR equation to control the relationship between adherence and the true SBP reduction from treatment. For example, to model a constant marginal benefit of treatment consumption, a uniform distribution function, $U(0, M)$ would be used for $\mathcal{F}_{\tau,s}$. In this example, the uniform distribution is not being used to claim that all SBP reductions between 0 and M are equally likely, but rather that there is a constant marginal benefit of improving treatment consumption. When combined with the CVaR equation, a uniform $\mathcal{F}_{\tau,s}$ indicates that each increase in treatment consumption α causes the same increase in SBP reduction, e.g. a 5% increase in adherence leads to a 1 mmHg higher reduction in SBP regardless of the patients initial adherence. Since CVaR provides the true SBP reduction from treatment, we add an error term to the CVaR SBP reduction to capture process and observation noise in our numerical studies. This error term captures the natural fluctuations observed in practice and the inaccuracy of blood pressure measurements.

We can map the patient's state s to the one-period pretreatment risk of a CHD event $r(s)$. Paired with the CVaR SBP reduction, we can then compute a relative risk factor $D(s) \in [0, 1]$ and resultant post-treatment risk $r_{\alpha,\tau}(s)$, according to [88]:

$$D(s) = \Delta(s)^{\frac{x_{\alpha,\tau}(s)}{20}}, \quad \Delta(s) \in (0, 1) \tag{3.5}$$

$$r_{\alpha,\tau}(s) = D(s) \cdot r(s) \tag{3.6}$$

The 20 in the equation for $D(s)$ represents a reference standard SBP reduction of 20 mmHg for which $\Delta(s)$ was computed. Note that since $x_{\alpha,\tau}(s)$ is increasing in α, then $D(s)$ is decreasing in α. Therefore, post-treatment CHD risk decreases as the patients treatment consumption increases.

Based on this post-treatment risk, the patient may enter one of four health states: (1) no CHD event; (2) survival of a CHD event; (3) death from a CHD event; or (4) death from a non-CHD cause. By construction, the patient has at most one event each decision period. Therefore, we assume the patient can have at most N CHD events in a planning horizon of N decision epochs. A surviving patient (i.e. health state 1 or 2) continues to the next decision epoch and repeats the process of selecting a treatment.

The patients pretreatment SBP, HDL and TC evolve over time as a function of the patients age. The dynamics of SBP, HDL and TC are captured using independent linear mixed effects (LME) models. We further assume those dynamics are affected by process noise that is added to the LME models. We also assume that hypertension treatment only affects the patients SBP and CHD risk in the period the treatment is taken. For example, if a patient discontinues treatment, the patients SBP would return to the pretreatment SBP forecasted by the LME model.

Our modeling framework provides the following set of dynamic programming equations:

$$V(s) = \min_{\gamma_\tau \leq y; \tau \in A(b_{min})} \left\{ \sum_{s' \in S} p_\tau(s, s')[J(s') + \lambda V(s')] \right\} \tag{3.7}$$

$$V(n = 0, \ldots) = 0 \tag{3.8}$$

Our recursive equations include a discount factor λ, but we set this parameter to 1 for our analysis.

The state transition probabilities are functions of the patients post-treatment CHD risk, $r_{\alpha,\tau}(s)$; the likelihood of death from a CHD event, $\rho(s)$; and the likelihood of non-CHD death, $\phi(s)$. Note that while not directly included in the notation, we adjust the likelihood of non-CHD death $\phi(s)$ to account for the patient's post-treatment CHD risk $r_{\alpha,\tau}(s)$. We first compute the post-treatment CHD risk $r_{\alpha,\tau}(s)$, which provides the total probability of transitioning to either health state 2 or 3. Next, we estimate $\phi(s)$ given $r_{\alpha,\tau}(s)$, denoted as simply $\phi(s)$. This ensures that $1 - r_{\alpha,\tau}(s) - \phi(s)$ (the transition probability to health state 1) is not negative and that the sum of all state transition probabilities equals 1. In solving the dynamic programming equations, we also manually check the transition probabilities to guarantee a valid transition probability distribution.

By definition, $J(\ldots, h = 2) = J(\ldots, h = 3) = 1$ (i.e. 1 CHD event occurred) and $J(\ldots, h = 1) = J(\ldots, h = 4) = 0$ (i.e. no CHD event occurred). In order to penalize death from a CHD event earlier in life, we set $V(n, \ldots, h = 3) = n$, i.e. the optimal value function is set to n events if the patient dies from a CHD event when there are n decision epochs remaining. And we set $V(n, \ldots, h = 4) = 0$ as there is no need to penalize for non-CHD deaths.

Combining the above, we can look at the optimal value function for health state 1 (i.e. no CHD event) as:

(3.9)

$$V(n,\ldots,h=1) = \min_{\gamma_\tau \leq y;\tau \in A(b_{min})} \{(1 - r_{\alpha,\tau}(s) - \phi(s))V(n-1,\ldots,h=1)$$

(3.10)
$$+ r_{\alpha,\tau}(s)(1-\rho(s))(1+V(n-1,\ldots,h=2))$$

(3.11)
$$+ r_{\alpha,\tau}(s)\rho(s)(1+V(n-1,\ldots,h=3))\}$$

Once a patient has had a CHD event (i.e. $h = 2$), the pretreatment CHD risk is multiplied by $(1+\zeta)$ to reflect the increased risk of subsequent CHD events due to a prior event. We assume only the first CHD event increases the patients CHD risk. We assume experiencing multiple CHD events does not further increase the subsequent CHD risk. In all living health states ($h = 1$ or $h = 2$) after the first CHD event, the patients CHD risk is scaled by $(1 + \zeta)$. Thus, when $h = 2$, in the computation for the next decision epoch optimal value function, we set $V(n-1,\ldots,h=1) = V(n-1,\ldots,h=2)$ to capture the continued increased CHD risk. For $h = 2$, we obtain the following optimal value function:

(3.12)

$$V(n,\ldots,h=2) = \min_{\gamma_\tau \leq y;\tau \in A(b_{min})} \{(1 - (1+\zeta)r_{\alpha,\tau}(s) - \phi(s))V(n-1,\ldots,h=1)$$

(3.13)
$$+ (1+\zeta)r_{\alpha,\tau}(s)(1-\rho(s))(1+V(n-1,\ldots,h=2))$$

(3.14)
$$+ (1+\zeta)r_{\alpha,\tau}(s)\rho(s)(1+V(n-1,\ldots,h=3))\}$$

(3.15)

If $\zeta > 0$ then $V(n,\ldots,h=1) < V(n,\ldots,h=2)$. Once again, we guarantee a valid transition probability distribution by computing $\phi(s)$ given $r_{\alpha,\tau}(s)$. We first enforce the requirement that $(1+\zeta)r_{\alpha,\tau}(s)$ is less than or equal to 1. Then we compute

the likelihood of a non-CHD death given the patients post-treatment CHD risk. A discount factor $\lambda \in (0, 1)$ may be considered as well, as is commonly done in the economic evaluation of health care interventions [36]. Since our model formulation for computing the expected weighted number of CHD events over the planning horizon penalizes more heavily death from a CHD event earlier in life and accounts for the increased likelihood of subsequent CHD events, we believe further discounting through inclusion of λ in the dynamic programming equations is unnecessary.

3.4 Model Properties

We now study structural properties of the optimal value function as they relate to treatment consumption heterogeneity and resource utilization. These properties are important to understanding current clinical practice and they provide insight into how guidelines can be altered to improve patient health. All proofs appear in the Section 3.10.1. We start by showing the intuitive result that as resource utilization increases, the minimized expected weighted number of CHD events is nonincreasing, i.e.

Theorem 1. $V(s)$ is nonincreasing with K.

This result suggests that not imposing restrictions on total treatment intensity for treating hypertension would benefit patients. Similarly, decreasing treatment intensity results in worse health outcomes. We now show that, in a high resource utilization environment, a myopic policy which selects the treatment that maximizes SBP reduction in that decision epoch (without violating the minimum allowable SBP constraint) is optimal, i.e.

Theorem 2. Let high resource utilization be defined as $y \geq n \max_{\tau \in A_n(b_{min};y)}\{\gamma_\tau\}$ when there are n decisions remaining. Under high resource utilization and if $x_{\alpha,\tau^*}(s) \geq$

$x_{\alpha,\tau}(s) \quad \forall \tau \in A_n(b_{min}; y)$, then τ^* is the optimal treatment when there are n decisions remaining.

The above result has several implications for clinical practice which we present as corollaries. First, under a mild assumption of nonincreasing minimum allowable SBP levels, the theorem implies that treatment intensity under the optimal policy never decreases over time. Then the MDP has a control-limit type policy under high resource utilization, which is a desirably structural property in MDP models [139].

Corollary 1. Let the treatments be ordered by intensity, i.e. $x_{\alpha,\tau_i}(s) < x_{\alpha,\tau_i+1}(s)$ for $i = 1, \ldots, |A|$. Under high resource utilization and $A_n(b_{min}; y) \subseteq A_k(b_{min}; y)$ for $n > k$, then τ^* is nondecreasing with decreasing n.

The theorem also implies that the optimal treatment policy will yield a SBP that approaches the minimum allowable SBP. If we consider the minimum allowable SBP to be the target SBP set by the physician, then the resource sufficiency theorem states that a target-based strategy is optimal.

Corollary 2. Under high resource utilization, the optimal policy leads to a post-treatment SBP that reaches (or approaches) the minimum allowable SBP, b_{min}. Thus, the optimal policy is TTT under resource sufficiency.

Next, we consider the implications of low resource utilization. In this scenario, the optimal policy is not guaranteed to be myopic, i.e. minimizing CHD risk in each epoch does not necessarily minimize the expected number of CHD events. With low resource utilization, it may be favorable to reserve resources (e.g. choose a less intense treatment in the current epoch) for later epochs when CHD risk is higher and treatments are more effective due to increasing SBP with age. This implies that with low resource utilization, the TTT strategy is not guaranteed to be optimal.

Corollary 3. Let low resource utilization be defined as $y < n \max_{\tau \in A_n(b_{min}; y)}\{\gamma_\tau\}$

when there are n decisions remaining. Under low resource utilization, $x_{\alpha,\tau^*}(s) \geq x_{\alpha,\tau}(s)$ does not guarantee τ^* is optimal.

While the TTT policy is optimal under high resource utilization, current U.S. guidelines have led to only 53.1% of patients with controlled blood pressure. In addition to low adherence, another reason for this poor control is the under-treatment of hypertension [42] and the restrictions imposed in clinical practice regarding increases in the intensity of prescribed medications. In Section 3.6, we evaluate the impact of those restrictions via our comparison of current U.S. guidelines against the MDP optimal policy which does not restrict the change in prescription intensity.

The poor control is also the result of the guidelines' incorporating neither risk factor heterogeneity nor the treatment consumption heterogeneity parameter of patient adherence to treatment. Our MDP approach explicitly captures the effects of heterogeneity on treatment outcomes to tailor the optimal treatment decisions. By incorporating heterogeneity, the MDP is capable of properly selecting the treatment intensity that will achieve the target SBP. Furthermore, knowledge of the patient's particular adherence level is necessary to prevent violations of the minimum allowable SBP constraint.

We now consider the effects of the treatment consumption heterogeneity parameter on the optimal value function and optimal treatment policy. First, we show the intuitive result of a nonincreasing minimized expected number of CHD events with an increasing treatment consumption heterogeneity parameter. As α increases (e.g. patient adherence increases), SBP reduction increases and the resulting expected number of CHD events decreases. In clinical practice, this theorem implies that improving patient adherence will yield improved health outcomes. In Section 3.6, we compute the marginal benefit of adherence on patient health outcomes. We

also analyze the tradeoff between an optimal policy which directly models patient adherence and policies which assume the patient is perfectly adherent to treatment.

Theorem 3. When the feasible action space $A_n(b_{min}; y)$ is the same across all α for each n, $V(s)$ is nonincreasing with α.

Finally, we examine how the treatment consumption heterogeneity parameter influences the decision of when to initiate treatment. Since treatment initiation decisions are a special case of the treatment intensity decisions, the timing of treatment initiation is equivalent to identifying the first decision epoch at which resources are expended.

Theorem 4. Let $\eta(\alpha)$ be the set of decision periods where it is optimal to not treat the patient, given the patient's treatment consumption heterogeneity parameter α. Then for $\bar{\alpha} > \alpha$, if:

1. The feasible action space $A_n(b_{min}; y)$ is the same under α and $\bar{\alpha}$ for each n,
2. $V_\alpha(s) - V_{\bar{\alpha}}(s) = c \geq 0$ for $\alpha \leq \bar{\alpha}$, i.e. equal change in the optimal value functions when moving between two α levels across all patient states.
3. $V(\ldots, h = 1) \leq (1 - \rho(s))V(\ldots, h = 2)$, i.e. the optimal value function when the patient has had a CHD event in the past is sufficiently larger than the optimal value function when the patient has not had a prior CHD event.

then $\eta(\bar{\alpha}) \subseteq \eta(\alpha)$.

The above implies that if a patient with a particular α chooses to treat in a given decision epoch, the same patient with a higher α will also choose to treat. Because the theorem is true for any number of remaining decision epochs, it is true for the first decision in the planning horizon. Therefore, a patient with a given α will begin treatment no later than a patient with identicalcharacteristics except for a smaller

α.

The first assumption prevents instances where the higher benefit of a treatment at higher α levels leads to infeasibility due to the minimum SBP constraint. Numerical studies have shown that the second assumption is reasonable, especially for problem instances with low resource utilization. And the third assumption is fair given the large increase in the risk of subsequent CHD events following a patient's first CHD event [124].

3.5 Identification of Heterogeneity Parameters

When tailoring treatment guidelines to patient adherence, it is important to have a reliable estimate of the patient's true adherence. A reliable estimate also allows a health provider to compute the marginal benefit of adherence on patient health outcomes. If the treatment consumption heterogeneity parameter is unknown, we can use the patient's history of SBP observations and treatments to dynamically estimate adherence. However, the identification of all elements α_{ij} in the matrix of parameters Θ is challenging. With assumptions on the structure of Θ, we can develop a method for updating our knowledge of the treatment consumption heterogeneity parameters over time.

Let the elements α_{ij} of Θ have a functional dependence f_{ij} on a single heterogeneity parameter α, i.e.

$$\Theta = \begin{bmatrix} f_{1,1}(\alpha) & \cdots & f_{1,N}(\alpha) \\ f_{2,1}(\alpha) & \cdots & f_{1,N}(\alpha) \\ \vdots & \ddots & \vdots \\ f_{|A|,1}(\alpha) & \cdots & f_{|A|,N}(\alpha) \end{bmatrix}$$

Then the identification and updating of α leads to the identification and updating

of the full matrix Θ. Furthermore, by construction, any nonstationarity for α is captured in the functions f_{ij}. In our analysis, we let $\Theta = \alpha\mathbf{J}$, where $\mathbf{J}$ is a $|A| \times N$ matrix of ones. Other forms for Θ are possible, given the necessary data. To identify and update our estimate of α, we start by defining $\hat{x}_{\alpha,\tau}(n)$ as the observed SBP reduction from the treatment selected when there are n decision epochs remaining and the true treatment consumption is α. This observed SBP reduction may be affected by noise, i.e.

$$\hat{x}_{\alpha,\tau}(n) = x_{\alpha,\tau}(n) + \epsilon \tag{3.16}$$

where $x_{\alpha,\tau}(n)$ is the true CVaR SBP reduction and ϵ is a noise term, e.g. $\epsilon \sim N(0, \sigma_\epsilon^2)$. Let $\alpha(\hat{x}_{\alpha,\tau}(n)) \in [0,1]$ be the treatment consumption which corresponds to the observed SBP reduction, calculated as:

$$\alpha(\hat{x}_{\alpha,\tau}(n)) = \sup_{\alpha\in[0,1]} \{\alpha : \hat{x}_{\alpha,\tau}(n) = CVaR_\alpha(X)]\} \tag{3.17}$$

If the treatment outcome distributions, $\mathcal{F}_{\tau,s}$, are 1-1 distribution functions, then $\alpha(\hat{x}_{\alpha,\tau}(n))$ is unique. When there is no closed form for equation (3.17), computational methods such as bisection [143] can be used to determine the corresponding treatment consumption heterogeneity parameter. By the stationarity of α, we can compute our estimate, $\hat{\alpha}_n$, of the true heterogeneity parameter α when there are n decision epochs remaining as the mean of all the sequentially observed corresponding heterogeneity parameters. With an initial estimate $\hat{\alpha}_N$ (which can be estimated using population data), we update our estimate via:

$$\hat{\alpha}_n = \frac{(N-n)\hat{\alpha}_{n+1} + \alpha(\hat{x}_{\alpha,\tau}(n+1))}{N-n+1} \tag{3.18}$$

3.6 Numerical Analysis

3.6.1 Parameterization

We utilized longitudinal data from the U.S. Department of Veterans Affairs, a popular CHD risk calculator [39], a meta-analysis of hypertension randomized clinical trials [88], and U.S. life tables [10] to parameterize our MDP model. Our VA data set combined data collected between 1999 and 2009 from the Veterans Health Administration (VHA) Medical SAS Data Sets, VHA Decision Support System files, and VHA Corporate Data Warehouse. The combined data contained information on approximately 2.5 million patients treated at VA hospitals across the U.S. For every patient in the data set, we had information on patient demographics (e.g. age, sex, diabetes) as well as longitudinal clinical observations (SBP, TC, HDL) and any prescribed medications. The average (standard deviation) follow-up time for patients in the data set was 6.4 (2.4) years. Table 1 summarizes the data sources and values for various model inputs.

From this VA data set, we randomly sampled 50,000 patients to use for parameterizing statistical models of stochastic risk factor evolution over time and to establish distributions for initial patient characteristics. Diabetes, alocholism, and smoking status were not available in the data. We assigned diabetes, alcoholism and smoking status to these patients at random based on their observed frequency in the U.S. population [1], [111], [27]. We used linear mixed effects (LME) models [121] to forecast the stochastic risk factors of SBP, TC, and HDL. LME models are statistical models which account for the inherent intra-patient correlation of longitudinal observations. Using the LME models and patient demographic data, we generated a sample patient population comprised of 3,000 patients followed for 10 years for use in the numerical study of our MDP model.

Parameter	Value	Source
Initial age	66.43 (11.05)	Analysis of VA Data
Percentage of men	97%	Analysis of VA Data
Percentage with diabetes	8.3%	[1]
Percentage of smokers	20%	[27]
Percentage of alcoholics	10%	[111]
Initial pretreatment SBP	142.12 (22.17) mmHg	Analysis of VA Data
Initial HDL	42.69 (13.08) mg/dL	Analysis of VA Data
Initial TC	180.16 (40.46) mg/dL	Analysis of VA Data
Pretreatment CHD risk, $r(s)$	Baseline 1-year risk: 0.0074 (0.0065); Baseline 10-year risk: 0.13 (0.10)	Analysis of VA Data & [39]
Minimum allowable SBP, b_{min}	120 mmHg	[14]
Treatment set, A	0 to 5 drugs, each at either half or full dosage	[88]
SBP reduction function, $F_\tau(s)$	$N(\mu_\tau, \sigma^2)$ with coefficient of variation = 1	[88]
Relative risk factor, $D(s)$	[0.49, 0.67]	[88]
Probability of death from non CHD cause, $\phi(s)$	[0.002, 0.08]	[10]
Probability of death from a CHD event given that the patient has a CHD event, $\rho(s)$	[0.095, 0.235]	[10]
Pretreatment CHD risk multiplier for subsequent CHD events, ζ	2	[124]

Table 3.1: Model inputs and data sources for hypertension treatment model.

To compute the likelihood of a CHD event in an epoch given the patient's risk factors, we used the risk calculator developed by researchers studying the Framingham Heart Study [39]. For each patient, we used the risk calculator to compute a pre-treatment CHD risk over the length of the decision epoch (e.g. 1 year CHD risk). The risk of CHD events after the patient had his/her first CHD event was computed using a scaling factor ζ, i.e. subsequent CHD event risk is $(1 + \zeta)$ times higher than the patient's risk for his/her first CHD event. Research shows that patients are at higher risk (estimated around 3 times higher) for subsequent CHD events than for initial CHD events, i.e. $\zeta = 2$ [124].

Analyses of randomized clinical trials revealed that SBP reduction from treatment depends on the number of prescribed drugs ($\eta \in \{0, 1, 2, 3, 4, 5\}$) and the dosage

level ($d \in \{\text{full}, \text{half}\}$) of the drugs [88]. We sorted (lowest to highest) and labeled the action space based on the intensity of treatments, which corresponds to the expected SBP reduction of the treatment. With this ranking, the lowest intensity treatment (intensity 1) was no treatment, intensity 2 was one half-standard dosage drug, intensity 3 is one full dosage drug, intensity 4 is two half-standard dosage drugs, intensity 5 is one full and one half-standard dosage drug, and so on with the highest intensity treatment (intensity 21) at five full dosage drugs.We used equations developed by [88] to compute the expected SBP reduction from a combination of drug count and dosage (d^{η}). We selected the Gaussian distribution, truncated at 0 mmHg reduction in SBP, as the outcome distribution for a particular treatment, where the mean is the computed expected SBP reduction and the variance is σ^2. We utilized a coefficient of variation (CV) of 1 to determine the variance. By construction, a patient with heterogeneity parameter $\alpha = 1$, would have a CVaR SBP reduction $x_{\alpha,\tau}$ equal to the expected SBP reduction. Estimates of relative CHD risk reduction due to SBP reduction, $\Delta(s)$ are also taken from [88]. Using the patient's pre-treatment CHD risk and the relative CHD risk reduction due to treatment, we computed the patient's post-treatment CHD risk which yielded our health state transition probabilities.

3.6.2 Results

With a parameterized MDP model, we first analyzed numerical results from the perspective of the treatment guideline policymaker. Using simulation, we compared the optimal treatment policy derived by the MDP against the TTT strategy for treating patients with hypertension (JNC7). We considered two implementations of the JNC7 guidelines: conservative and aggressive. Both implementations initialized with the same treatments (either one or two drugs per the guidelines). The conservative JNC7 implementation increased one half-dosage drug to full dosage or added one

half-dosage drug if all currently prescribed medications were already at full dosage when the patient had not reached the target SBP. The aggressive JNC7 implementation prescribed one additional full dosage drug to the patient when he/she had not reached the target SBP. In our simulation, each patient was treated annually over 10 years according to both the conservative and aggressive JNC7 implementation. First, a patient was simulated under the JNC7 implementations and the total number of prescribed medications (i.e. total treatment intensity) over the 10 year horizon was computed for conservative and aggressive JNC7. The total treatment intensity was then used as the resource input K in the MDP model (solved under the K from conservative JNC7 and under the K from aggressive JNC7). This ensured that the JNC7 guidelines and MDP treatment policies expended equal amounts of resources on any given patient.

Furthermore, we considered levels of the treatment consumption heterogeneity parameter for adherence ranging between $\alpha = 0.1$ and $\alpha = 1$. Since JNC7 does not directly incorporate adherence in the treatment guidelines, implementations of JNC7 guidelines assumed the patient's true adherence was $\alpha = 1$ at all times. For the MDP model, we initialized with $\alpha = 1$ for all patients for a more fair comparison with JNC7, but used our identification algorithm from Section 3.5 to learn the true α over time given noisy clinical observations of SBP. After each clinical observation of SBP reduction (the true SBP reduction plus Gaussian white noise ϵ, considered at high and low variance), our estimate of α was updated according to equation (3.18) and the MDP was solved for the remaining decision periods.

We tested the sample population of 3,000 patients at each level of the heterogeneity parameter, each level of observation noise, and each JNC7 implementation and MDP model. Each patient was simulated 500 times. The simulation model uses

the same inputs and parameters described in Section 3.6.1 and Table 3.1. For each scenario, we computed the following performance measures: the average number of CHD events and the average time until the patient's first CHD event.

Figure 3.3 illustrates the relative improvement in the average number of CHD events of the MDP policy over the JNC7 guidelines. The figure presents the results for the four combinations of JNC7 implementation (conservative and aggressive) and observation noise (low and high) at various levels of the treatment consumption heterogeneity parameter. The average percentage reduction in the number of CHD events over the 10 year planning horizon ranged from 6.5% to 9.4% ($p < 0.05$ for all improvements using paired t-tests). In addition to examining differences for the mean number of CHD events, we investigated differences within the distribution of patients. Patients with high initial SBP, i.e. > 140 mmHg, achieved the greatest reduction in the average number of CHD events (18.7% and 18.0%, compared against conservative JNC7 at $\alpha = 1$ with low and high noise, respectively). When compared against aggressive JNC7, the relative improvement for high initial SBP patients decreased (11.7% and 9.8% at $\alpha = 1$ with low and high noise, respectively), but the reduction still exceeded the overall average improvement.

While time until the first event (TTFE) was not the objective of the MDP model, Figures 3.4 and 3.5 shows that the MDP policies resulted in later first CHD events for patients who had CHD events ($p < 0.05$ for all MDP and JNC7 comparisons using paired t-tests). This performance measure is considered clinically important, because a patient's first CHD event increases the likelihood of subsequent CHD events during the patient's lifetime. The average percentage improvement in the TTFE over the 10 year planning horizon ranged from 5.1% to 6.6%. Similar to the relative improvement in the number of CHD events, patients with high initial SBP achieved the greatest

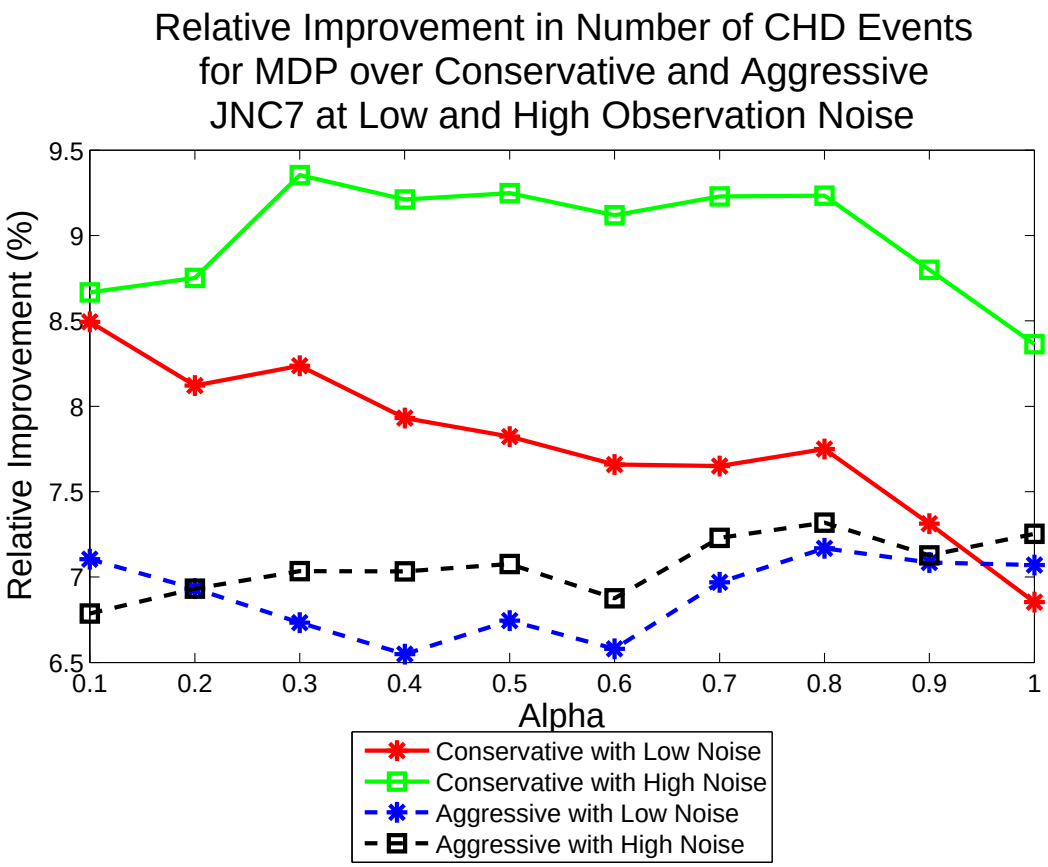

Figure 3.3: Percentage improvements in the average number of CHD events.

improvement in TTFE (18.5% and 17.0%, compared against conservative JNC7 at $\alpha = 1$ with low and high noise, respectively). Compared against aggressive JNC7, the relative improvement for high initial SBP patients decreased (15.5% and 10.4% at $\alpha = 1$ with low and high noise, respectively), but the improvement still exceeded overall average improvement in TTFE. Furthermore, this average 3 to 4 month delay in TTFE per patient corresponds to a significant increase in quality of life for the population as a whole. With 785,000 new coronary events in the United States every year [128], the delays in first heart attacks from the MDP treatment policy corresponds to approximately 31,400 quality-adjusted life years saved [120, 119].

In general, we found that the relative improvement in health outcomes was higher

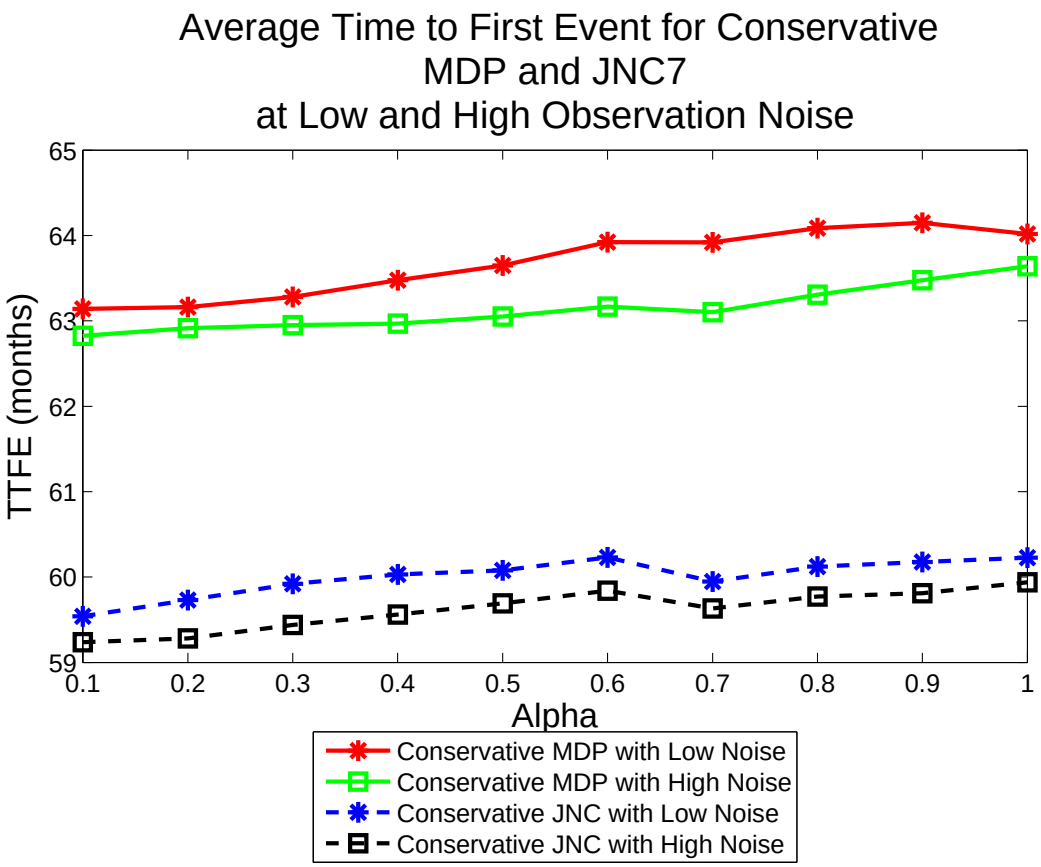

Figure 3.4: Conservative JNC7 implementation

when comparing the MDP policy against the conservative JNC7 implementation. This result is due to ceiling effects for the aggressive JNC7, i.e. the aggressive implementation yields higher intensity treatment decisions which correspond to increased likelihoods of reaching the target SBP. Thus, the improvements achievable through optimization by the MDP policy over the aggressive implementation are relatively smaller than the improvements achievable over the conservative implementation.

In addition to using simulation to compare the MDP policy against the TTT strategy for a sample population of patients, we investigated the marginal benefit of resources (total treatment intensity) on the number of CHD events. Figure 3.6 plots the expected number of CHD events over the next 10 years for a sample

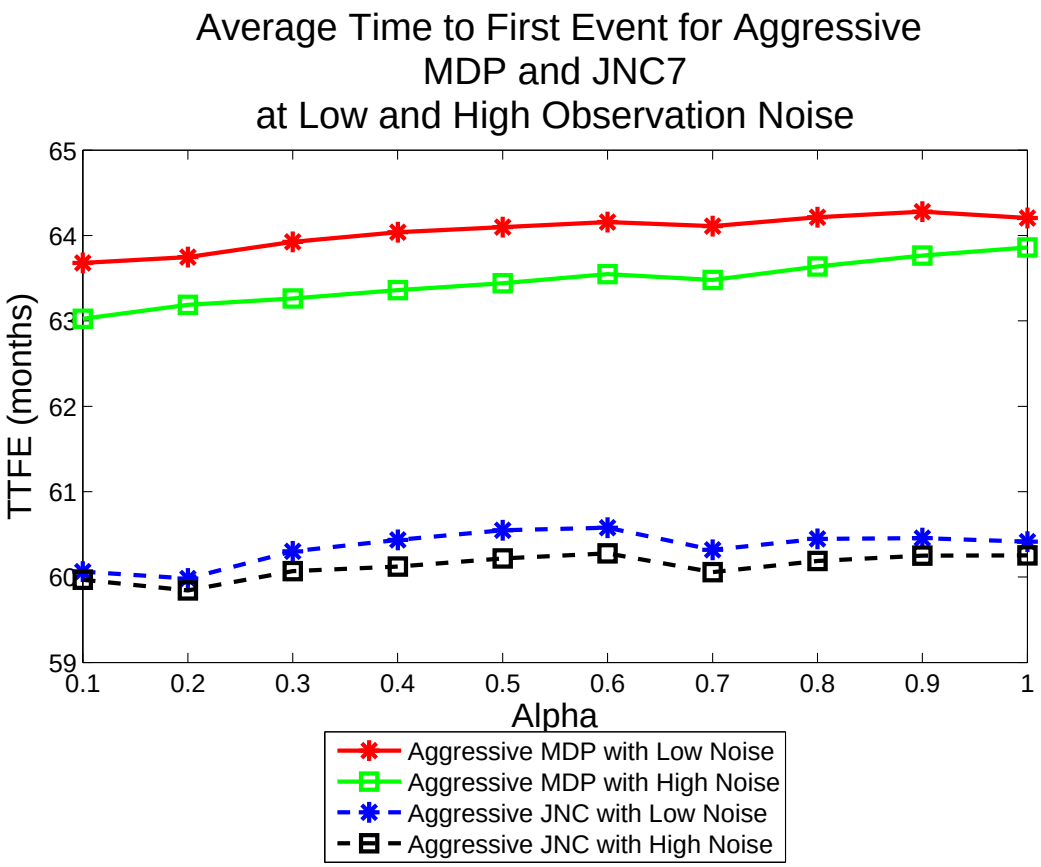

Figure 3.5: Aggressive JNC7 implementation

patient at various total treatment intensities. The figure shows that the number of CHD events decreases with increasing resources (Theorem 1). By our inclusion of the resource utilization in the state space, the resource curve plotted in Figure 3.6 can be generated immediately after solving the MDP. These resource curves are patient-specific and allow for easy cost-effectiveness analyses. From the figure, we can see regions of different marginal improvements in the objective function. There are regions of large marginal improvements, small marginal improvements, and no marginal improvements. These resource curves can help clinicians and treatment guideline policymakers interpret the value of intensifying treatment intensity on the patient's expected weighted number of CHD events over his/her lifetime.

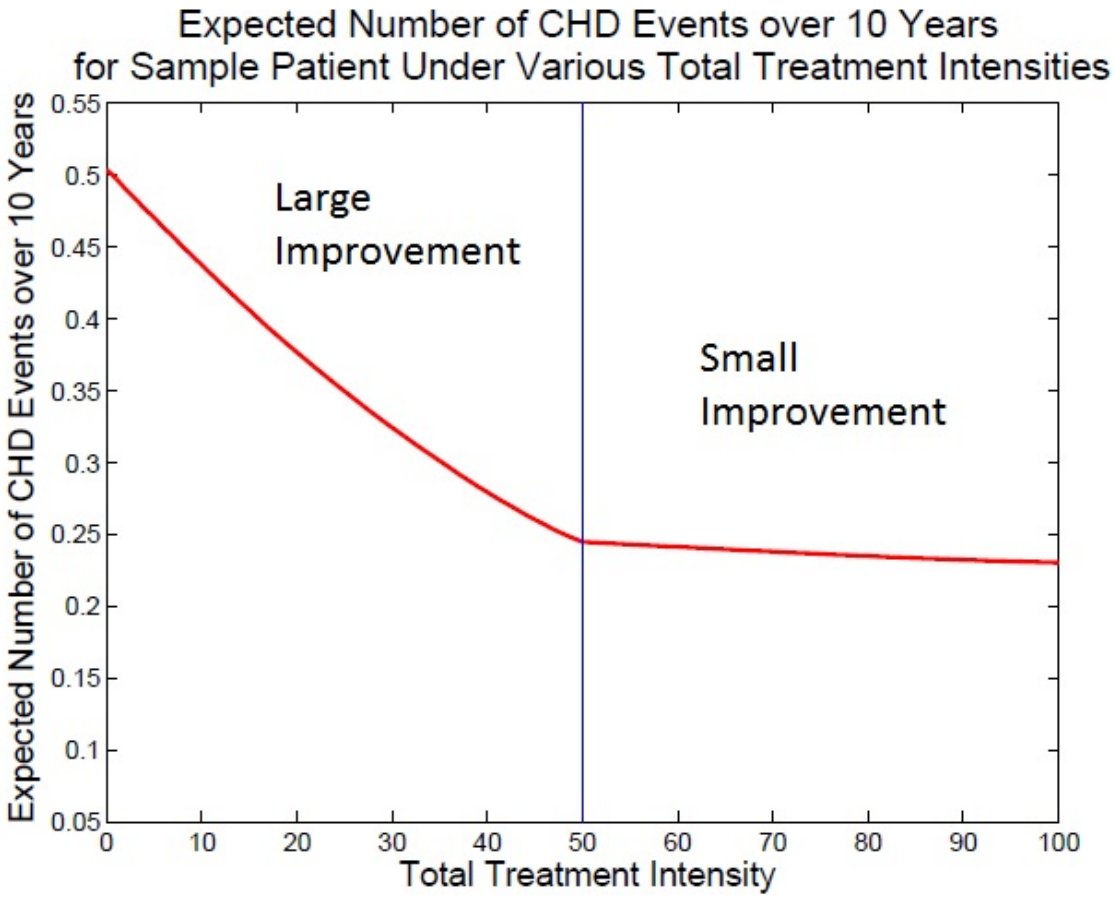

Figure 3.6: Comparison of the expected number of CHD events at various total treatment intensities for a sample patient.

For the same patient, we can also visualize the intensity of treatments in the MDP policy over the planning horizon. Figure 3.7 shows the optimal treatment policy for a sample patient when the total treatment intensity is high. We see that with resource utilization, the treatment intensity for optimal policy does not decrease over time (Corollary 1). Instead, the optimal treatment policy selects the highest intensity treatment that does not violate the minimum allowable SBP constraint (Theorem 2). Since the underlying pretreatment SBP is increasing with age, the highest intensity and feasible treatment is also increasing to bring the post-treatment SBP to the target (Corollary 2). However, the nondecreasing property of the treatment intensity

no longer holds when we consider the same patient under lower resource utilization (Corollary 3). Figure 3.8 shows that under low total treatment intensity the optimal policy starts at a milder intensity, increases briefly, then decreases toward the end of the planning horizon as resources are consumed.

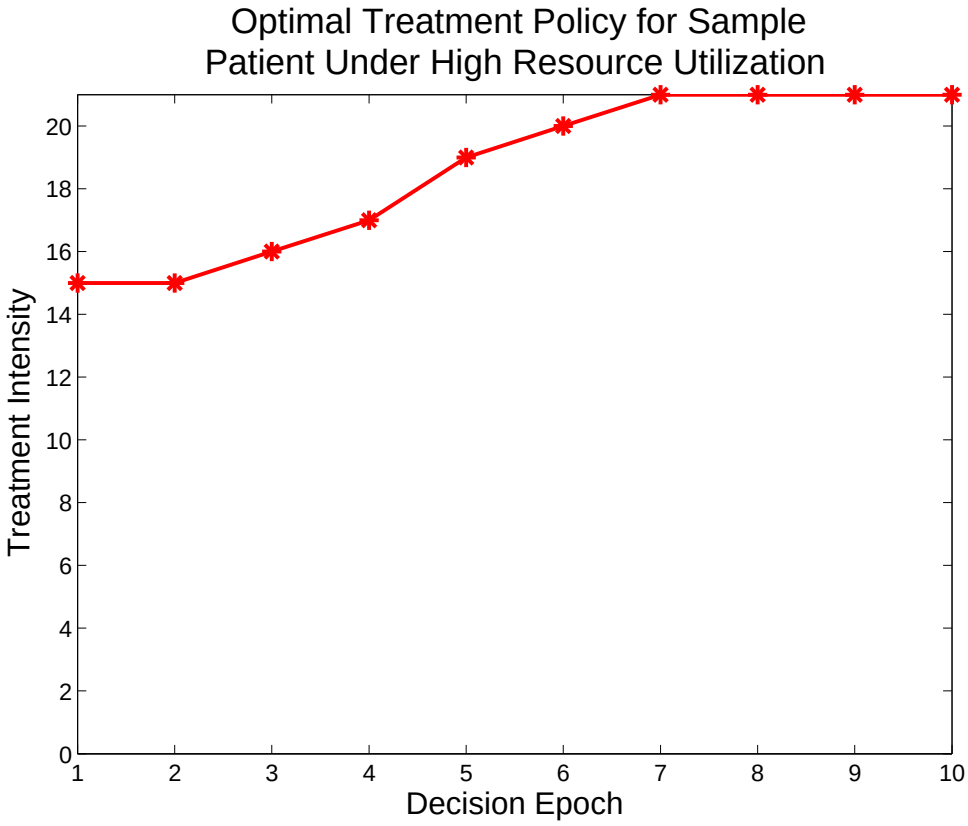

Figure 3.7: Optimal treatment policy under high resource utilization

Next, we considered two possible implementations of the optimal treatment policy within the sample population: (1) treatment policies which directly account for the patients treatment consumption heterogeneity and (2) treatment policies which assume all patients are fully adherent to treatment (i.e. $\alpha = 1$). Note that both implementations account for risk factor heterogeneity when making treatment decisions. The first implementation requires tracking the patients adherence level over time to tailor medication choices to his/her particular adherence. However, such an

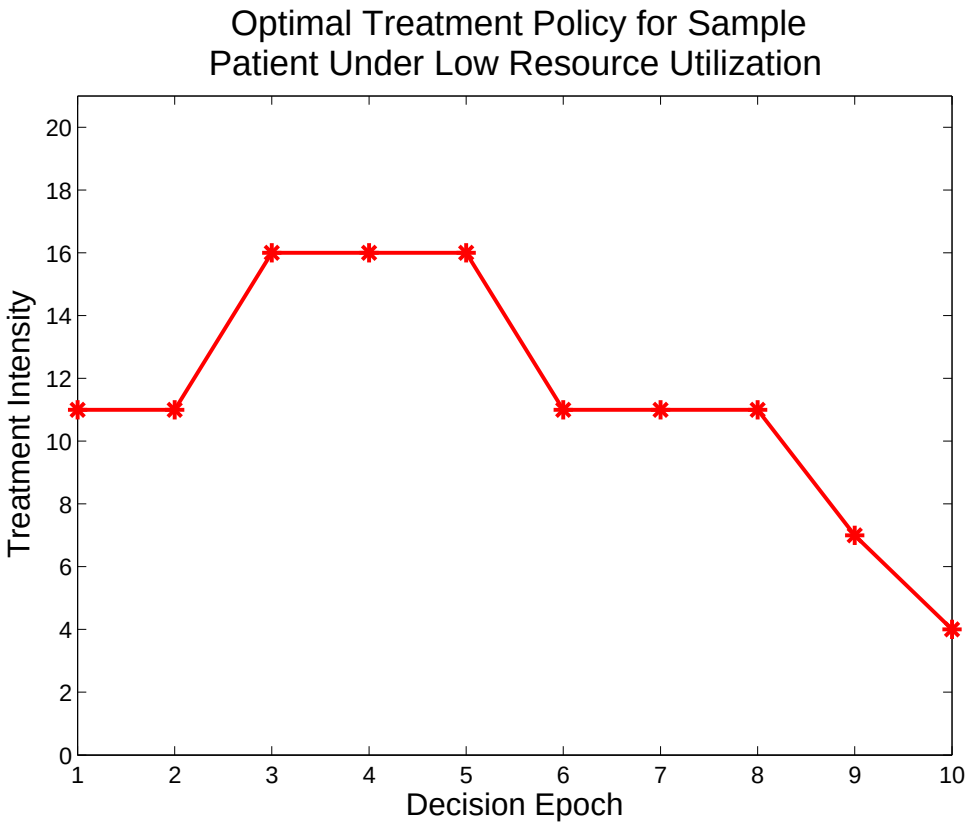

Figure 3.8: Optimal treatment policy under low resource utilization

implementation may be difficult to use in clinical practice. The second implementation mitigates those challenges by assuming that all patients will fully adhere to prescribed medications. Figure 3.9 shows the total number of CHD events over the next 10 years under these two policies per 1000 patients at true adherence levels between 0.1 and 1. Treatment policies which account for the patients true adherence outperform treatment policies which assume full adherence. Treatment guideline policymakers must decide whether the improvement in health outcomes is worth the additional complexity in decision making and implementation.

Finally, we analyzed numerical results from the perspective of the health provider or insurance company interested in evaluating adherence improving interventions. The marginal benefit of improving the adherence for the sample patient population is

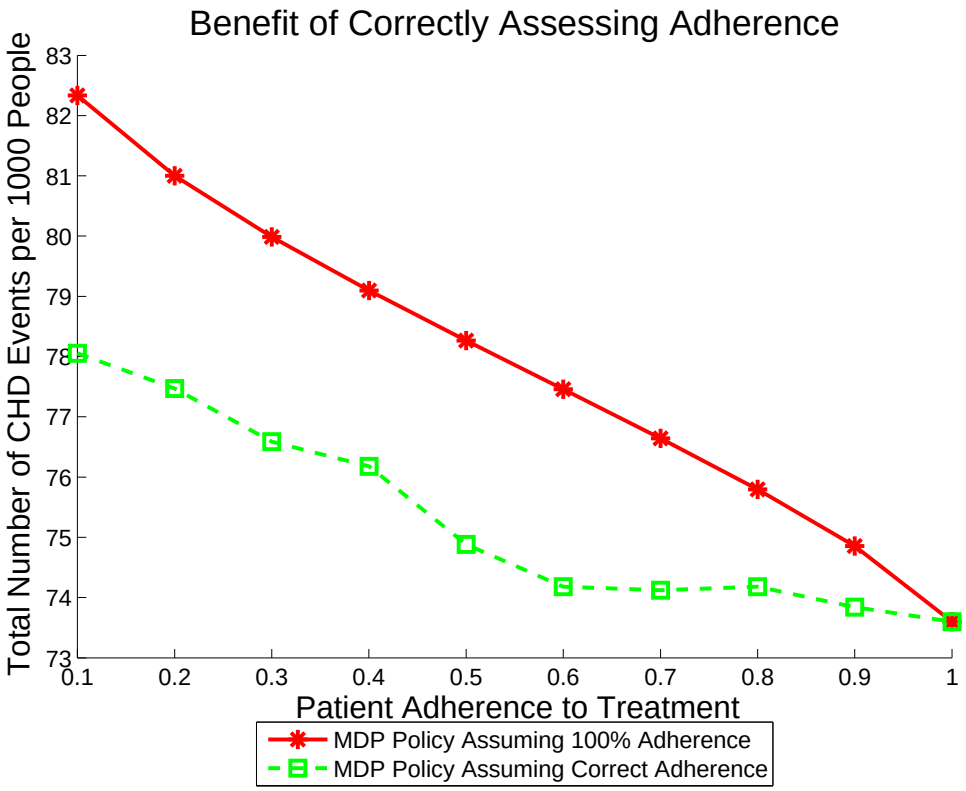

Figure 3.9: Comparison of CHD events per 1000 ppl under two implementations of the optimal treatment policy derived by the MDP.

shown in Figure 3.9. Paired with a cost of adherence interventions, cost-effectiveness analysis can be readily computed under each of the treatment policies: (1) the optimal policy which assumes all patients are fully adherent, or (2) the optimal policy which directly accounts for each patients particular adherence. Figure 3.9 also indicates that treatment guideline policymakers and health providers can work together to improve patient health outcomes. For example, if treatment guideline policymakers advocate for optimal policies which account for risk factor heterogeneity while assuming all patients will fully adhere to the prescribed medications, health providers and insurance companies can implement adherence improving interventions so that the assumption of full adherence is more appropriate. Our numerical analysis indicates that jointly addressing problems in the clinical management of hypertension

patients will generate the greatest gains in patient health.

3.7 Discussion

We have developed a modeling framework that utilizes the power of CVaR for the identification and exploitation of patient heterogeneity in stochastic sequential decision making. CVaR is an efficient and general approach for incorporating heterogeneity in treatment outcomes in the derivation of optimal treatment policies. We applied this modeling framework to the problem of lifetime hypertension treatment decisions for patients at risk of CHD events and studied the structural properties of the MDP formulation. We focused on three possible clinical management policy changes: (1) using optimal treatment policies which incorporate patient heterogeneity, (2) increasing resource utilization when treating patients, and (3) utilizing interventions which improve treatment consumption. Our analysis revealed that under high resource utilization, a control-limit type policy is optimal when knowledge of the heterogeneity parameters, such as adherence to therapy, is used to decide the treatment intensity. In addition to numerically verifying the structural properties, we utilized a large VA dataset of longitudinal clinical observations to parameterize our MDP and simulation models which compared the TTT strategy for treating hypertension against the treatment policies generated by our MDP model. Our optimal MDP policies yielded fewer CHD events and later first CHD events for patients when compared to the performance of the TTT strategy. Our analysis also reveals that by working jointly, treatment guideline policymakers and health providers can achieve maximum gains in patient health through optimal treatment planning and the implementation of adherence improving interventions.

We note that new recommendations have been released for the management of

hypertension [44]. Since these recommendations have not been incorporated into clinical practice, we focused on comparisons against the widely used JNC7 guidelines. The recommendations (target SBP of 150 mm Hg) are considerably less aggressive (i.e. lower resource utilization) in treating hypertension than the conservative JNC7 implementation tested in this paper. We expect that our MDP framework would achieve greater benefits than the improvements numerically shown over the JNC7 guidelines.

Further research remains to be done in calibrating and implementing the model discussed. While our work has focused on adherence as a summary variable of treatment consumption heterogeneity, other elements of heterogeneity in medical decision making (e.g. biological response) can be included in the dynamic programming formulation to improve its real-world performance. Currently, our modeling framework makes use of a single matrix of heterogeneity parameters for treatment consumption, Θ. Incorporating multiple sources of treatment consumption heterogeneity separately (e.g. letting α denote education and β denote socioeconomic status) would require a joint distribution F_τ so that CVaR is computed given α and β. Using this joint distribution, we could then evaluate the impact of these heterogeneity parameters on treatment outcomes. While a joint distribution may be easily modeled, such a formulation creates additional complexity when inverting the distribution to learn the treatment consumption parameters given SBP reductions observed over time. Similarly, while we assumed in the numerical analysis that adherence was constant across treatments and age, data pertaining to the factors which influence adherence could be used to enrich the functional form of the heterogeneity matrix. Also, our assumption that hypertension treatment only affects the patients SBP and CHD risk in the period the treatment was prescribed may be unrealistic. For example, the benefit of

treatment may depend on the number of decision periods the patient has been taking medication. We leave the analysis of such treatment phenomena to future research. We used the Framingham risk calculator, but other calculators with different risk factors may prove to be better estimators of CHD event likelihoods. Further understanding of the impact of distribution assumptions remains, and other distributions (e.g. exponential and uniform) for the SBP reduction from treatments need to be considered. It would also be beneficial to examine the health benefits of our MDP model when applied to developing nations that operate under resource insufficiency. Lastly, while our application area was hypertension treatments for patients at risk of CHD, our modeling framework is generalizable to other chronic diseases, particularly those where severely adverse events result from disease progression.

3.8 Model Extension: Optimizing Quality-Adjusted Life Years

To extend the research described above, we next consider treatment policies which maximize the patient's expected discounted quality-adjusted life years (QALYs). We use this new model to answer two clinically relevant questions: (1) how does misestimation of treatment disutility impact treatment decisions and patient outcomes and (2) what is the impact of time-varying treatment benefit and disutility on current treatment strategies. To increase the acceptance of our results by the clinical community, we include both CHD and stroke as adverse health outcomes. The inclusion of stroke requires extending our state space to now contain 10 mutually-exclusive patient health states: (1) healthy (no history of CHD or stroke); (2) history of CHD but no CHD event this period; (3) history of stroke but no stroke this period; (4) history of CHD and stroke but no adverse event this period; (5) survived a CHD event this period; (6) survived a stroke this period; (7) death from a non-CVD related cause;

(8) death from CHD event this period; (9) death from stroke this period; and (10) dead. We also expand the dataset upon which the analysis is performed by utilizing data from the National Health and Nutrition Examination Survey (NHANES), which contains data on a representative sample of the U.S. population [148]. Furthermore, for this analysis, we assume all patients are fully adherent (i.e. $\alpha = 1$), and we leave the study of how adherence would impact these models for future research.

The QALY-maximizing MDP is characterized by the following recursive dynamic programming equations for the optimal value function $V(s)$:

$$V(s) = \max_{\tau \in A(b_{min})} \left\{ \sum_{s' \in S} p_\tau(s, s')[J(s', \tau) + \lambda V(s')] \right\} \tag{3.19}$$

where $J(s', \tau)$ equals the quality of life associated with being in state s' minus the treatment disutility d_τ of treatment τ. In addition to the immediate reward's dependence on the treatment, equation (**??**) differs from equation (3.7) in that there is no resource constraint $\gamma_\tau \leq y$. Furthermore, in the original dynamic program, the terminal condition was set to 0 (i.e. no CHD events after the planning horizon ends). However, in this analysis of QALYs, we compute a terminal condition as the product of the patients expected lifetime [9], a mortality scaling factor [21], and the quality of life weighting for the health state. All data inputs and sources for the QALY-maximizing MDP can be found in Section 3.10.2.

3.8.1 Physician Estimation of Treatment Disutility

One reason for low adherence to medication is the disutility from treatment, specifically the side effects and costs. All medications have side effects and costs to both the patient and society as a whole. However, the measurement of treatment disutility varies between patients [50], [71], and it is challenging to estimate in a clinical visit. If clinicians incorrectly estimate the disutility a patient places on a medication, the

clinician may not prescribe the best treatment plan to the patient. As differences between a physician's perception of treatment disutility and the patient's true disutility increase [166], the rate of mistreatment is likely to increase. Clinicians with the belief that treatments have low disutility may overtreat patients who believe the disutility is high. Similarly, clinicians who believe treatments have low disutility may undertreat patients who believe the disutility is low.

To compute the impact of the physician's ability to accurately estimate a patient's true treatment disutility, we estimate the differences in the optimal treatment plan and the expected outcomes of that plan when the patient's true disutility and the physician's perceived disutility are misaligned. We assume at the beginning of the planning horizon that all patients have never had a CHD event or stroke (i.e. the first health state). We assume medication dosage decisions are based upon the physician's perceived disutility of treatment after solving equation (3.19) for the QALY-maximizing MDP. The resulting MDP treatment policy is then evaluated under the patient's true treatment disutility. We consider four levels of disutility per antihypertensive medication: none (0), low (0.001), medium (0.005), and high (0.001). Each patient in our sample population was evaluated under the MDP for every combination of true and perceived disutility. We then repeated this process using the JNC7 guidelines which are not sensitive to the perceived disutility. For each scenario (disutility combination and treatment policy), we computed the expected total number of QALYs, the expected number of total medications prescribed, the expected total number of CV events, and the total number of patients treated. We compared the performance of the MDP against the current treatment guidelines of JNC7. We also compared the MDP policy under the true disutility against the MDP policies for each of the perceived disutilities to detect significant differences in

treatments (i.e. mistreatment rates).

Table 3.2 summarizes the performance measures as applied to 100,000 randomly generated patients from NHANES III followed over 10 years. Five performance measures were recorded: the expected number of patients treated, the expected total number of medications prescribed per 1000 treated patients, the expected number of CV events prevented per 1000 patients, the expected change in QALYs per 1000 patients, and the expected change in QALYs per 1000 treated patients. To be classified as a treated patient, the patient needed to have been prescribed at least 5 medications over the 10 year planning horizon. Performance measures on QALYs and number of CV events were computed relative to JNC7. Notice that, for each perceived disutility, the optimal treatment policy is the same for a single patient across all true treatment disutilities. Therefore, there are no changes in the expected number of treated patients, number of medications prescribed, or number of CV events within a perceived disutility.

As one may expect, as the perceived disutility of treatment increases, the number of patients treated decreases. Patients at low risk for CV events are unlikely to receive enough benefit from treatment to outweigh the disutility and hence would never be prescribed hypertension medication. Similarly, the total number of medications prescribed amongst those who are treated also decreases as the perceived treatment disutility increases. However, as the total number of medications prescribed decreases, overall CV risk is higher which leads to fewer CV events prevented compared to the current guidelines of JNC7. The greater number of CV events causes fewer QALYs in the population. All of these findings are true regardless of the patient's true disutility.

Table 3.3 presents the mistreatment rate (percentage of patients with a difference

		Perceived Disutility			
Performance Measure	**True Disutility**	0	0.001	0.005	0.01
Expected number of patients	0	59171	59109	54116	46562
treated	0.001	59171	59109	54116	46562
	0.005	59171	59109	54116	46562
	0.01	59171	59109	54116	46562
Expected total number of meds	0	24670.11	24634.72	23287.43	21511.4
prescribed per 1000 ppl Rx'd	0.001	24670.11	24634.72	23287.43	21511.4
	0.005	24670.11	24634.72	23287.43	21511.4
	0.01	24670.11	24634.72	23287.43	21511.4
Expected number of CV events	0	17.62	17.62	16.23	12.83
prevented per 1000 ppl	0.001	17.62	17.62	16.23	12.83
	0.005	17.62838	17.62	16.23	12.83
	0.01	17.62838	17.62	16.23	12.83
Expected QALYs saved	0	111.93	111.90	106.10	89.18
per 1000 ppl	0.001	105.16	105.17	101.13	86.51
	0.005	78.07	78.21	81.22	75.80
	0.01	44.21	44.52	56.34	62.41
Expected QALYs saved per	0	1433.26	1430.48	1348.70	933.94
1000 ppl Rx'd	0.001	1426.63	1423.89	1343.52	930.58
	0.005	1400.11	1397.52	1322.79	917.17
	0.01	1366.97	1364.56	1296.89	900.39

Table 3.2: Performance of the QALY maximizing MDP under each combination of true and perceived treatment disutility when compared against JNC7.

of at least 0.5 medications per year between the MDP policy under the perceived disutility and the MDP policy under the true disutility) for each combination of true and perceived disutility. When the perceived disutility is greater than the true disutility, the mistreatment rate indicates the percentage of patients undertreated, whereas when the perceived disutility is less than the true, the mistreatment indicates overtreatment. As the difference between true and perceived disutility increases, the mistreatment rate also increases. The highest mistreatment rate occurs at the combination of no disutility and high disutility with 26.9% of patients being mistreated under the MDP policy due to differences between patient and physician calculations of treatment disutility.

We have found that the patient's perception of treatment disutility should play a key role in the hypertension medication decision-making process. When patients are either over- or under-treated due to differences between the patient's true disutility

		Perceived Disutility			
	% Mistreated	0	0.001	0.005	0.01
True Disutility	0	0.0%	0.2%	13.0%	26.9%
	0.001	0.2%	0.0%	13.0%	26.9%
	0.005	13.0%	13.0%	0.0%	19.8%
	0.01	26.91%	26.9%	19.8%	0.0%

Table 3.3: Percentage of patients mistreated for each combination of true and perceived disutility.

and the physician's perceived disutility, there is a loss of QALYs. While the number of CV events is always minimized when the perceived disutility is 0, the costs of medication outweigh the benefit of prevented CV events when true and perceived disutilities are misaligned. This misalignment of disutilities results in the mismanagement of patients, particularly when considering no disutility and high disutility. In addition to mistreatment, the misalignment of disutility may lead to poor adherence rates to prescribed medication, a motivating factor in the research presented in Section 3.3. Future research may address the need to effectively model the impact of disutility misestimation on patient adherence when desigining treatment policies.

Our model provides computations of the costs associated with improper estimation of a patient's true treatment disutility. We illustrate these costs through different levels of assumed disutility; unfortunately, identification of the patient's true disutility remains a challenge. Because patient disutility is inherently subjective, there are few tools to accurately estimate a given patient's views on medication trade-offs. However, our analysis reveals the importance of accurately estimating disutility and provides strong evidence for further research in this under-studied area.

This research also shows the short-comings of the current clinical guidelines of JNC7. Even with high misestimation of disutility, the MDP policies greatly outperform the JNC7 guidelines. The MDP formulation enables the treatment policy to consider the medication costs and trade-offs. In combination with a focus on CV risk rather than intermediate markers, such as SBP, the dynamic programming approach

to treatment planning offers superior health outcomes for patients.

3.8.2 Time-Varying Treatment Benefit and Disutility

Next, we use the data sources and Markov model discussed in 3.8 to evaluate the impact of time-varying treatment benefit and disutility on different treatment initiation strategies. In this context, the relative risk reduction (RRR) from treatment $D(s)$, used to compute $J(s, \tau)$, and the treatment disutility d_τ depend on how long the patient has been taking the medication. Our prior work in CV treatment planning assumed that the RRR (i.e. treatment benefit) was invariant to the length of time the patient had been taking the medication. Similarly for disutility, we assumed that the disutility per medication was constant across time. In this extension, we model the treatment benefit and disutility as a function of the patient's number of years on medication by expanding the state space to include the medication exposure duration η_τ, and we study how different assumptions on these functional relationships affect the performance of clinician-designed treatment strategies. In this research, we consider a 45 year planning horizon where patients may initiate treatment as early as age 30, and we follow them through age 75.

We analyzed four functional forms for the RRR from treatment. The first case assumes that the RRR is 30% all years the patient is on treatment. The second through fourth cases assume that the RRR is 30% for the first 5 years and then increases linearly to become 60% in year 10, 20, 30 (respectively) and remains 60% (i.e. maximum treatment benefit) in all future years. Figure 3.10 illustrates how the RRR depends on the number of years the patient has been on treatment.

We also considered five functional forms for the disutility from treatment. Figure 3.11 illustrates how treatment disutility depends on the number of years the patient has been on treatment. The first three cases assume a static 0.001, 0.005, and 0.01

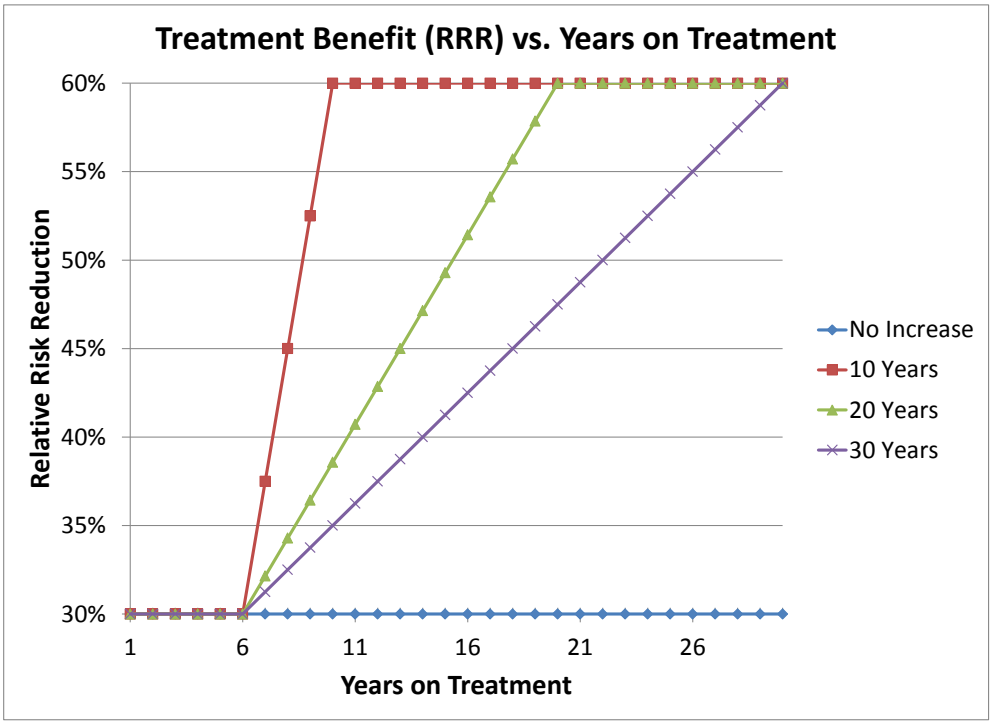

Figure 3.10: Relative risk reduction (RRR) from treatment as a function of the number of years the patient has been on treatment.

disutility. The last two cases initiate with a disutility of 0.001 in the first year of treatment and increase to 0.05 and 0.1, respectively, in year 25, according to the following equations:

$$d_\tau(\eta_\tau) = d_\tau(1+i)^{\eta_\tau} \tag{3.20}$$

$$i = 1 - \left(\frac{d_\tau^M}{d_\tau^m}\right)^{\frac{1}{25}} \tag{3.21}$$

where d_τ^m and d_τ^M denote the treatment disutility at year 1 and year 25, respectively, and $d_\tau(\eta_\tau)$ is the treatment disutility when the patient has been on treatment for η_τ years.

We considered treatment strategies, derived by our clinical collaborators, which

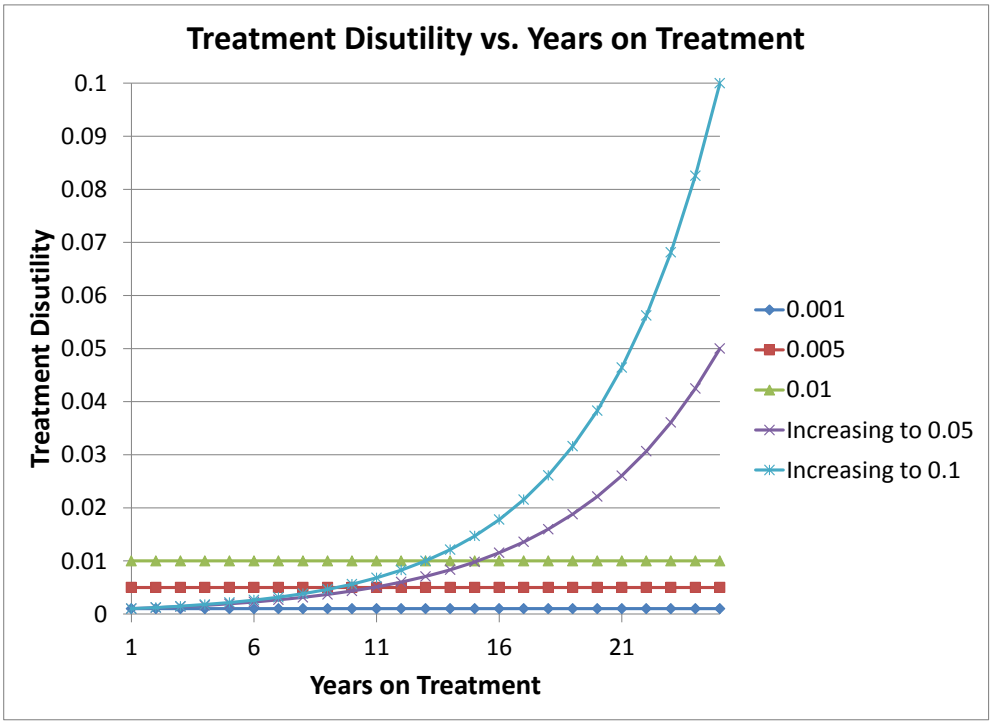

Figure 3.11: Treatment disutility per medication as a function of the number of years the patient has been on treatment.

initiate treatment when the patients current 10-year cardiovascular (CV) risk, equal to the sum of 10-year CHD and stroke risk as calculated by [8], exceeds a given threshold. By using strategies designed by clinicians, our analysis may be easier to implement in real-world practice. We considered thresholds of 5, 10, 15, and 20% CV risk. We also considered treatment strategies which initiate treatment when the patients forecasted 10-year CV risk 10 years in the future exceeds the given threshold. This future risk strategy initiates treatment early so the treatment benefit is closer to maximum benefit at the time the patients CV risk is high.

We applied each treatment strategy (current and future 10-year CV risk thresholds) to each patient beginning at age 30 for each combination of treatment benefit

and disutility function. We assumed patients never discontinue treatment. For each patient, treatment strategy and assumed treatment benefit and disutility function, we computed the expected number of treatment years, the expected number of CV events, and the expected QALYs. We then stratified the sample population based upon their age-adjusted CV risk (i.e. CV risk at age 50) into quartiles to analyze how the treatment strategies differ across risk stratums.

3.8.3 Results

Tables 3.4 through 3.6 present the results of different assumptions on treatment benefit and disutility for those patients in the top quartile of age-adjusted CV risk under each of the treatment strategies. For the treatment disutilities, 0.05^* and 0.1^* denote the increasing disutilities. Section 3.10.3 contains the analysis of the other three quartiles.

Comparing the assumption of no increase in RRR to the assumption of 10 years until full benefit is achieved, we found that the future 5% 10-year CV risk strategy is most affected by time-varying benefit. For an assumed low disutility (0.001), the expected QALYs per 1000 patients for the future 5% CV risk strategy under no increase in RRR was 23,435 and 24,089 when treatment takes 10 years to reach full benefit. This change in QALYs (654 QALYs per 1000 patients) was the highest amongst all treatment strategies indicating that the health outcome of the strategy highly depends on the modeling assumptions. Furthermore, for the low disutility, the future 5% CV risk strategy also attained the highest QALYs per 1000 patients amongst the tested strategies for all assumptions on time-varying treatment benefit. At higher disutilities, including the increasing disutility functions, no single strategy dominates the other strategies in terms of QALYs across all time-varying treatment benefit assumptions. In general, the future 5% CV risk strategy attained the highest

Top Quartile Age-Adjusted CV Risk Patients	**Rx years per 1000 ppl**				
Treatment Disutility	0.001	0.005	0.01	0.05*	0.1*
Current risk strategy					
No increase in RRR					
5% CV risk	34082	34082	34082	34082	34082
10% CV risk	27259	27259	27259	27259	27259
15% CV risk	21786	21786	21786	21786	21786
20% CV risk	17158	17158	17158	17158	17158
10 years till full benefit					
5% CV risk	35112	35112	35112	35112	35112
10% CV risk	28185	28185	28185	28185	28185
15% CV risk	22590	22590	22590	22590	22590
20% CV risk	17847	17847	17847	17847	17847
20 years till full benefit					
5% CV risk	35010	35010	35010	35010	35010
10% CV risk	28050	28050	28050	28050	28050
15% CV risk	22450	22450	22450	22450	22450
20% CV risk	17727	17727	17727	17727	17727
30 years till full benefit					
5% CV risk	34859	34859	34859	34859	34859
10% CV risk	27894	27894	27894	27894	27894
15% CV risk	22316	22316	22316	22316	22316
20% CV risk	17623	17623	17623	17623	17623
Future risk strategy					
No increase in RRR					
5% CV risk	39157	39157	39157	39157	39157
10% CV risk	36185	36185	36185	36185	36185
15% CV risk	31189	31189	31189	31189	31189
20% CV risk	26253	26253	26253	26253	26253
10 years till full benefit					
5% CV risk	40228	40228	40228	40228	40228
10% CV risk	37239	37239	37239	37239	37239
15% CV risk	32186	32186	32186	32186	32186
20% CV risk	27162	27162	27162	27162	27162
20 years till full benefit					
5% CV risk	40158	40158	40158	40158	40158
10% CV risk	37152	37152	37152	37152	37152
15% CV risk	32071	32071	32071	32071	32071
20% CV risk	27027	27027	27027	27027	27027
30 years till full benefit					
5% CV risk	40026	40026	40026	40026	40026
10% CV risk	37009	37009	37009	37009	37009
15% CV risk	31915	31915	31915	31915	31915
20% CV risk	26877	26877	26877	26877	26877

Table 3.4: Treatment (Rx) years for the top quartile age-adjusted CV risk patients under different assumptions on time-varying treatment benefit and disutility.

QALYs compared to the other strategies. The high performance of the future 5% CV risk strategy for patients in the top quartile of CV risk is due to the increased

Top Quartile Age-Adjusted CV Risk Patients	**CV events per 1000 ppl**				
Treatment Disutility	0.001	0.005	0.01	0.05*	0.1*
Current risk strategy					
No increase in RRR					
5% CV risk	767.27	767.27	767.27	767.27	767.27
10% CV risk	781.93	781.93	781.93	781.93	781.93
15% CV risk	804.86	804.86	804.86	804.86	804.86
20% CV risk	833.98	833.98	833.98	833.98	833.98
10 years till full benefit					
5% CV risk	427.24	427.24	427.24	427.24	427.24
10% CV risk	466.90	466.90	466.90	466.90	466.90
15% CV risk	520.39	520.39	520.39	520.39	520.39
20% CV risk	581.73	581.73	581.73	581.73	581.73
20 years till full benefit					
5% CV risk	450.38	450.38	450.38	450.38	450.38
10% CV risk	499.72	499.72	499.72	499.72	499.72
15% CV risk	558.83	558.83	558.83	558.83	558.83
20% CV risk	621.08	621.08	621.08	621.08	621.08
30 years till full benefit					
5% CV risk	486.38	486.38	486.38	486.38	486.38
10% CV risk	542.49	542.49	542.49	542.49	542.49
15% CV risk	601.49	601.49	601.49	601.49	601.49
20% CV risk	656.94	656.94	656.94	656.94	656.94
Future risk strategy					
No increase in RRR					
5% CV risk	763.85	763.85	763.85	763.85	763.85
10% CV risk	765.38	765.38	765.38	765.38	765.38
15% CV risk	772.13	772.13	772.13	772.13	772.13
20% CV risk	785.43	785.43	785.43	785.43	785.43
10 years till full benefit					
5% CV risk	414.95	414.95	414.95	414.95	414.95
10% CV risk	420.58	420.58	420.58	420.58	420.58
15% CV risk	440.70	440.70	440.70	440.70	440.70
20% CV risk	474.81	474.81	474.81	474.81	474.81
20 years till full benefit					
5% CV risk	431.98	431.98	431.98	431.98	431.98
10% CV risk	440.83	440.83	440.83	440.83	440.83
15% CV risk	467.79	467.79	467.79	467.79	467.79
20% CV risk	508.37	508.37	508.37	508.37	508.37
30 years till full benefit					
5% CV risk	462.11	462.11	462.11	462.11	462.11
10% CV risk	474.33	474.33	474.33	474.33	474.33
15% CV risk	506.96	506.96	506.96	506.96	506.96
20% CV risk	551.00	551.00	551.00	551.00	551.00

Table 3.5: Cardiovascular (CV) disease events for the top quartile age-adjusted CV risk patients under different assumptions on time-varying treatment benefit and disutility.

exposure duration η_τ over the other strategies. By initiating treatment when the patient is predicted to have at least 5% 10-year CV risk in 10 years from now, the

Top Quartile Age-Adjusted CV Risk Patients	**QALYs per 1000 ppl**				
Treatment Disutility	0.001	0.005	0.01	0.05*	0.1*
Current risk strategy					
No increase in RRR					
5% CV risk	23429	23356	23266	23120	22838
10% CV risk	23393	23342	23278	23204	23035
15% CV risk	23337	23300	23253	23216	23111
20% CV risk	23268	23241	23207	23185	23113
10 years till full benefit					
5% CV risk	24054	23980	23888	23731	23435
10% CV risk	23950	23897	23832	23750	23572
15% CV risk	23817	23778	23731	23689	23577
20% CV risk	23674	23646	23611	23585	23507
20 years till full benefit					
5% CV risk	23992	23918	23826	23670	23375
10% CV risk	23871	23818	23753	23671	23493
15% CV risk	23733	23695	23648	23605	23492
20% CV risk	23596	23568	23533	23506	23428
30 years till full benefit					
5% CV risk	23907	23834	23742	23587	23294
10% CV risk	23780	23728	23663	23582	23406
15% CV risk	23650	23612	23565	23522	23411
20% CV risk	23529	23501	23467	23440	23362
Future risk strategy					
No increase in RRR					
5% CV risk	23435	23345	23231	23011	22619
10% CV risk	23433	23354	23254	23080	22755
15% CV risk	23418	23355	23276	23163	22933
20% CV risk	23385	23336	23275	23207	23049
10 years till full benefit					
5% CV risk	24089	23997	23882	23648	23239
10% CV risk	24074	23993	23892	23705	23365
15% CV risk	24020	23956	23876	23753	23510
20% CV risk	23930	23880	23818	23743	23576
20 years till full benefit					
5% CV risk	24044	23952	23837	23605	23196
10% CV risk	24020	23940	23838	23652	23313
15% CV risk	23952	23887	23807	23685	23442
20% CV risk	23851	23802	23740	23664	23497
30 years till full benefit					
5% CV risk	23970	23878	23763	23533	23127
10% CV risk	23940	23859	23758	23574	23237
15% CV risk	23863	23799	23720	23598	23358
20% CV risk	23762	23713	23651	23577	23411

Table 3.6: Quality-adjusted life years (QALYs) for the top quartile age-adjusted CV risk patients under different assumptions on time-varying treatment benefit and disutility.

strategy is treating patients earlier and for longer. As long as the assumed treatment disutility is relatively low, the increased exposure duration leads to higher QALYs

for these high risk patients regardless of how long it takes the medication to reach full benefit.

This research has presented a first step toward modeling time-varying treatment benefit and disutility in the evaluation of treatment strategies. We expanded the state space to include an exposure duration η_τ parameter. This parameter provides a tractable way to incorporate general functions relating treatment benefit and disutility to how long a patient has been on medication. These functional forms and the numerical analyses performed provide insight into the trade-off between early initiation for higher RRR and the increasing cumulative disutility from longer medication exposure. While we focused on testing the impact of time-varying benefit and disutility on clinician-designed treatment strategies, our modeling approach enables the derivation of the optimal initiation strategy. We can also expand the methodology to include the initiation of multiple medications over a patient's lifetime (e.g. start with a statin, then start an antihypertensive medication later in life) through the addition of extra η_τ for each new treatment added to the patient's medication regimen.

3.9 Conclusion

In conclusion, we have presented a dynamic programming formulation for determining optimal hypertension treatment plans. Our utilization of conditional value-at-risk for incorporating patient heterogeneity allows for modeling a rich set of assumptions about the relationship between adherence and treatment outcome when computing immediate rewards. We showed that by personalizing optimal treatment policies, we are able to improve patient health outcomes compared to current treatment guidelines. Furthermore, our modeling framework allows for direct computation

of the marginal benefit of improving patient adherence to medication. This marginal benefit analysis informs decisions made at the population level, as we will discuss in Chapter IV. In addition to deriving optimal treatment policies and analyzing the effects of patient heterogeneity on health outcomes, we used the dynamic programming models to answer clinical questions about the impact of differences between true and perceived treatment disutility as well as the effect of time-varying treatment benefit and disutility on treatment strategies.

3.10 Appendix

3.10.1 Proofs

Proof of Theorem 1 (Optimal Value Function and Resources): Let Π_K be the set of feasible policies when the lifetime resource amount available for expenditure is K. And let $\pi_K^* \in \Pi_K$ be the optimal policy, i.e. $\pi_K^* = \arg\min_{\pi_K \in \Pi_K} V_{\pi_K}(s)$. Then if $K' > K$, then $\Pi_K \subseteq \Pi_{K'}$ which implies $V_{\pi_{K'}^*}(s) \leq V_{\pi_K^*}(s) \forall s \in S$. ■

Proof of Theorem 2 (Myopic Policy): Noting that $V(n-1, y-\gamma_\tau, \ldots, h=4) = 0$, the state-action cost function q_τ for treatment τ is:

$$\begin{aligned} q_\tau(n, y, \ldots, h=1) = (1 - r_{\alpha,\tau}(s) - \phi(s))V(n-1, y-\gamma_\tau, \ldots, h=1) \\ + r_{\alpha,\tau}(s)(1-\rho(s))[1 + V(n-1, y-\gamma_\tau, \ldots, h=2)] \\ + r_{\alpha,\tau}(s)\rho(s)[1 + V(n-1, y-\gamma_\tau, \ldots, h=3)] \end{aligned}$$

Let τ^* be defined such that $x_{\alpha,\tau^*}(s) \geq x_{\alpha,\tau}(s) \quad \forall \tau \in A_n(b_{min}; y)$. Then $r_{\alpha,\tau^*}(s) \leq r_{\alpha,\tau}(s) \quad \forall \tau \in A_n(b_{min}; y)$. To show optimality, we require that the state-action cost function difference of τ^* and all other τ is less than or equal to 0. If $y \geq n \max_{\tau \in A(b_{min}; y)}\{\gamma_\tau\}$, then:

$$\begin{aligned} V(n-1, y-\gamma_{\tau^*}, \ldots, h=1) &= V(n-1, y-\gamma_\tau, \ldots, h=1) \\ &= V(n-1, \ldots, h=1) \quad \forall \tau \in A(b_{min}; y) \\ V(n-1, y-\gamma_{\tau^*}, \ldots, h=2) &= V(n-1, y-\gamma_\tau, \ldots, h=2) \\ &= V(n-1, \ldots, h=2) \quad \forall \tau \in A(b_{min}; y) \\ V(n-1, y-\gamma_{\tau^*}, \ldots, h=3) &= V(n-1, y-\gamma_\tau, \ldots, h=3) \\ &= V(n-1, \ldots, h=3) \quad \forall \tau \in A(b_{min}; y) \end{aligned}$$

Which implies:

$$\begin{aligned} q_{\tau^*} - q_\tau = (r_{\alpha,\tau^*}(s) - r_{\alpha,\tau}(s))[&-V(n-1, y-\gamma_\tau, \ldots, h=1) \\ &+ (1-\rho(s))(1+V(n-1, y-\gamma_\tau, \ldots, h=2)) \\ &+ \rho(s)(1+V(n-1, y-\gamma_\tau, \ldots, h=3))] \end{aligned}$$

By construction, $V(n-1, y-\gamma_\tau, \ldots, h=1) \leq V(n-1, y-\gamma_\tau, \ldots, h=2) \leq V(n-1, y-\gamma_\tau, \ldots, h=3) = n-1$ and $V \geq 0$. Since $r_{\alpha,\tau^*}(s) - r_{\alpha,\tau}(s) \leq 0$, we require for optimality of τ^* that:

$$\begin{aligned} (1-\rho(s))V(n-1, y-\gamma_\tau, \ldots, h=2) - \\ V(n-1, y-\gamma_\tau, \ldots, h=1) \geq -\rho(s)(n-1) - 1 \end{aligned}$$

The above is always true since $V(n-1, y-\gamma_\tau, \ldots, h=1) \leq V(n-1, y-\gamma_\tau, \ldots, h=2) \leq V(n-1, y-\gamma_\tau, \ldots, h=3) = n-1$. For all $\rho(s)$ values between 0 and 1, the requirement for optimality holds. Therefore, τ^* is optimal. ■

Proof of Corollary 1: By the above theorem, the optimal treatment with n decisions remaining when $y \geq n \max_{\tau \in A_n(b_{min}; y)}\{\gamma_\tau\}$ is $\tau^* : x_{\alpha,\tau}(s) \leq x_{\alpha,\tau^*}(s) \quad \forall \tau \in A_n(b_{min}; y)$. If $A_n(b_{min}; y) \subseteq A_k(b_{min}; y)$ for $n > k$, then τ^* is feasible in all future decision epochs. The optimal treatment changes only if $x_{\alpha,\tau^*}(s) < x_{\alpha,\tilde{\tau}}(s)$ for some $\tilde{\tau} \in A_k(b_{min}; y)$. By $x_{\alpha,\tau_i}(s) < x_{\alpha,\tau_i+1}(s)$ for $i = 1, \ldots, |A|$, this implies τ^* is nondecreasing with decreasing n. ■

Proof of Corollary 2: By the above theorem, the optimal treatment with n decisions remaining when $y \geq n \max_{\tau \in A_n(b_{min}; y)}\{\gamma_\tau\}$ is $\tau^* : x_{\alpha,\tau}(s) \leq x_{\alpha,\tau^*}(s) \quad \forall \tau \in A_n(b_{min}; y)$. The set of feasible treatments $A_n(b_{min}; y)$ is constrained by the minimum allowable SBP, b_{min}, such that feasible treatments do not lead to a SBP below b_{min}. Therefore, the optimal treatment when there are n decisions remaining minimizes the

distance between the post-treatment SBP, $b_n - x_{\alpha,\tau}(s)$, and the minimum allowable SBP, b_{min}, i.e. $\tau^* = \arg\min_{\tau \in A_n(b_{min};y)}(b_n - x_{\alpha,\tau}(s)) - b_{min}$. ■

Proof of Corollary 3: If $y < n \max_{\tau \in A_n(b_{min};y)}\{\gamma_\tau\}$, and $x_{\alpha,\tau^*}(s) \geq x_{\alpha,\tau}(s)$ then $\gamma_{\tau^*} > \gamma_\tau$. By the nonincreasing with resources property of the optimal value function: $V(n-1, y-\gamma_{\tau^*}, \ldots, h=1) \geq V(n-1, y-\gamma_\tau, \ldots, h=1)$. Thus, $q_{\tau^*}(n, y, \ldots, h = 1) - q_\tau(n, y, \ldots, h = 1)$ is not guaranteed to be ≤ 0, as it was in the proof under the assumption of high resource utilization. Therefore, the myopic policy is not guaranteed to be optimal. ■

Proof of Theorem 3 (Optimal Value Function and Heterogeneity Parameter): Let $\alpha^* > \alpha$. Then $r_{\alpha^*,\tau}(s) < r_{\alpha,\tau}(s)$ because $x_{\alpha^*,\tau}(s) > x_{\alpha,\tau}(s)$. For a fixed y, τ and n:

$$\begin{aligned}
q_{\alpha^*} - q_\alpha = {} & (1 - r_{\alpha^*,\tau}(s) - \phi(s))V_{\alpha^*}(n-1, y-\gamma_\tau, \ldots, h=1) \\
& - (1 - r_{\alpha,\tau}(s) - \phi(s))V_\alpha(n-1, y-\gamma_\tau, \ldots, h=1) \\
& + r_{\alpha^*,\tau}(s)(1-\rho(s))[1 + V_{\alpha^*}(n-1, y-\gamma_\tau, \ldots, h=2)] \\
& - r_{\alpha,\tau}(s)(1-\rho(s))[1 + V_\alpha(n-1, y-\gamma_\tau, \ldots, h=2)] \\
& + r_{\alpha^*,\tau}(s)\rho(s)[1 + V_{\alpha^*}(n-1, y-\gamma_\tau, \ldots, h=3)] \\
& - r_{\alpha,\tau}(s)\rho(s)[1 + V_\alpha(n-1, y-\gamma_\tau, \ldots, h=3)]
\end{aligned}$$

Consider the final decision epoch, i.e. $n = 1$. The terminal condition is given as $V(n = 0, \ldots) = 0$. Therefore:

$$q_{\alpha^*}(n, y, \ldots, h=1) - q_\alpha(n, y, \ldots, h=1) = r_{\alpha^*,\tau}(s) - r_{\alpha,\tau}(s) < 0$$

Because the above is true for any τ, we conclude that $V_{\alpha^*}(n = 1, \ldots) \leq V_\alpha(n = 1, \ldots)$. The inequality is strict when the patient is in a living health state (i.e. $h = 1$ or $h = 2$).

We now have:

$$(r_{\alpha,\tau}(s)+\phi(s))V_{\alpha}(n-1,y-\gamma_{\tau},\ldots,h=1)-$$
$$(r_{\alpha^*,\tau}(s)+\phi(s))V_{\alpha^*}(n-1,y-\gamma_{\tau},\ldots,h=1)\geq 0$$

Since $r_{\alpha,\tau}(s)+\phi(s)\leq 1$ and $r_{\alpha^*,\tau}(s)+\phi(s)\leq 1$, we have the following relationships:

$$(r_{\alpha,\tau}(s)+\phi(s))V_{\alpha}(n-1,y-\gamma_{\tau},\ldots,h=1)\leq V_{\alpha}(n-1,y-\gamma_{\tau},\ldots,h=1)$$
$$(r_{\alpha^*,\tau}(s)+\phi(s))V_{\alpha^*}(n-1,y-\gamma_{\tau},\ldots,h=1)\leq V_{\alpha^*}(n-1,y-\gamma_{\tau},\ldots,h=1)$$

Therefore, we have:

$$\begin{aligned}V_{\alpha^*}(n-1,y-\gamma_{\tau},\ldots,h=1)-V_{\alpha}(n-1,y-\gamma_{\tau},\ldots,h=1)&\\ +(r_{\alpha,\tau}(s)+\phi(s))V_{\alpha}(n-1,y-\gamma_{\tau},\ldots,h=1)&\\ -(r_{\alpha^*,\tau}(s)+\phi(s))V_{\alpha^*}(n-1,y-\gamma_{\tau},\ldots,h=1)&\\ \leq 0&\end{aligned}$$

Then the sum of the inequalities is <0. Hence, $V_{\alpha^*}(n=2,\ldots)\leq V_{\alpha}(n=2,\ldots)$. By induction, this relationship holds for all n. ■

Proof of Theorem 4 (Initial Treatment Timing and Heterogeneity Parameter): Let $\eta(\alpha)=\{n\mid q_{\alpha,0}(n,y,\ldots,h=1)-q_{\alpha,\tau}(n,y,\ldots,h=1)\leq 0\}$, i.e. the set of decision periods where it is optimal to not treat the patient. To prove that $\eta(\bar{\alpha})\subseteq\eta(\alpha)$ for $\bar{\alpha}>\alpha$, it is necessary to prove that $q_{\alpha,0}-q_{\alpha,\tau}$ is increasing as α

increases.

$$\begin{aligned} q_{\alpha,0} - q_{\alpha,\tau} &= (1 - r(s) - \phi(s))V(n-1, y, \ldots, h=1) \\ &\quad + r(s)(1-\rho(s))[1 + V(n-1, h, \ldots, h=2)] \\ &\quad + r(s)\rho(s)[1 + (n-1)] \\ &\quad - (1 - r_{\alpha,\tau}(s) - \phi(s))V(n-1, y-\gamma_\tau, \ldots, h=1) \\ &\quad - r_{\alpha,\tau}(s)(1-\rho(s))[1 + V(n-1, y-\gamma_\tau, \ldots, h=2)] \\ &\quad - r_{\alpha,\tau}(s)\rho(s)[1 + (n-1)] \end{aligned}$$

Let $A_0(\alpha)$ and $A_\tau(\alpha)$ be defined as:

$$\begin{aligned} A_0(\alpha) &= (1 - r(s) - \phi(s))V(n-1, y, \ldots, h=1) \\ &\quad + r(s)(1-\rho(s))[1 + V(n-1, h, \ldots, h=2)] \\ A_\tau(\alpha) &= (1 - r_{\alpha,\tau}(s) - \phi(s))V(n-1, y-\gamma_\tau, \ldots, h=1) \\ &\quad + r_{\alpha,\tau}(s)(1-\rho(s))[1 + V(n-1, y-\gamma_\tau, \ldots, h=2)] \end{aligned}$$

Then we have:

$$q_{\alpha,0} - q_{\alpha,\tau} = A_0(\alpha) - A_\tau(\alpha) + n\rho(s)[r(s) - r_{\alpha,\tau}(s)]$$

Since $r_{\alpha,\tau}(s)$ decreases as α increases, then $n\rho(s)[r(s) - r_{\alpha,\tau}(s)]$ is increasing with α. It remains to be shown that $A_0(\alpha) - A_\tau(\alpha)$ is increasing with α. Let $Q(\alpha) =$

$A_0(\alpha) - A_\tau(\alpha)$. If $Q(\alpha) - Q(\bar{\alpha}) \leq 0$, then we have completed the proof.

$$
\begin{aligned}
Q(\alpha) - Q(\bar{\alpha}) &= A_0(\alpha) - A_\tau(\alpha) - (A_0(\bar{\alpha}) - A_\tau(\bar{\alpha})) \\
&= (1 - r(s) - \phi(s))V_\alpha(n-1, y, \ldots, h=1) \\
&+ r(s)(1-\rho(s))[1 + V_\alpha(n-1, h, \ldots, h=2)] \\
&- (1 - r_{\alpha,\tau}(s) - \phi(s))V_\alpha(n-1, y-\gamma_\tau, \ldots, h=1) \\
&- r_{\alpha,\tau}(s)(1-\rho(s))[1 + V_\alpha(n-1, y-\gamma_\tau, \ldots, h=2)] \\
&- (1 - r(s) - \phi(s))V_{\bar{\alpha}}(n-1, y, \ldots, h=1) \\
&- r(s)(1-\rho(s))[1 + V_{\bar{\alpha}}(n-1, h, \ldots, h=2)] \\
&+ (1 - r_{\bar{\alpha},\tau}(s) - \phi(s))V_{\bar{\alpha}}(n-1, y-\gamma_\tau, \ldots, h=1) \\
&+ r_{\bar{\alpha},\tau}(s)(1-\rho(s))[1 + V_{\bar{\alpha}}(n-1, y-\gamma_\tau, \ldots, h=2)]
\end{aligned}
$$

Under the assumption that $V_\alpha(s) - V_{\bar{\alpha}}(s) = c \geq 0$, then we have:

$$
\begin{aligned}
Q(\alpha) - Q(\bar{\alpha}) &= (1 - r(s) - \phi(s))c \\
&\quad + r(s)(1 - \rho(s))c \\
&\quad + (1 - \rho(s))[r_{\bar{\alpha},\tau}(s) - r_{\alpha,\tau}(s)] \\
&\quad - r_{\alpha,\tau}(s)(1 - \rho(s))V_\alpha(n-1, y-\gamma_\tau, \ldots, h=2) \\
&\quad + r_{\bar{\alpha},\tau}(s)(1 - \rho(s))V_{\bar{\alpha}}(n-1, y-\gamma_\tau, \ldots, h=2) \\
&\quad - V_\alpha(n-1, y-\gamma_\tau, \ldots, h=1) + V_{\bar{\alpha}}(n-1, y-\gamma_\tau, \ldots, h=1) \\
&\quad + r_{\alpha,\tau}(s)V_\alpha(n-1, y-\gamma_\tau, \ldots, h=1) \\
&\quad - r_{\bar{\alpha},\tau}(s)V_{\bar{\alpha}}(n-1, y-\gamma_\tau, \ldots, h=1) \\
&\quad + \phi(s)V_\alpha(n-1, y-\gamma_\tau, \ldots, h=1) - \phi(s)V_{\bar{\alpha}}(n-1, y-\gamma_\tau, \ldots, h=1) \\
Q(\alpha) - Q(\bar{\alpha}) &= -r(s)\rho(s)c + (1 - \rho(s))[r_{\bar{\alpha},\tau}(s) - r_{\alpha,\tau}] \\
&\quad + r_{\alpha,\tau}(s)[V_\alpha(n-1, y-\gamma_\tau, \ldots, h=1) \\
&\quad - (1 - \rho(s))V_\alpha(n-1, y-\gamma_\tau, \ldots, h=2)] \\
&\quad - r_{\bar{\alpha},\tau}(s)[V_{\bar{\alpha}}(n-1, y-\gamma_\tau, \ldots, h=1) \\
&\quad + (1 - \rho(s))V_{\bar{\alpha}}(n-1, y-\gamma_\tau, \ldots, h=2)]
\end{aligned}
$$

Assuming $V_\alpha(\ldots, h=1) \leq (1 - \rho(s))V_\alpha(\ldots, h=2)$, then $Q(\alpha) - Q(\bar{\alpha}) \leq 0$. ■

3.10.2 Data for QALY-Maximizing MDP

Table 3.7 summarizes the inputs and data sources for the model.

Table 3.8 presents the 10 mutually-exclusive health states and their respective quality of life weight.

Table 3.9 reports the mortality ratio and quality of life weights for each terminal health state.

Risk slopes used to compute relative risk reduction are shown in Table 3.10

Input	Description	Source
T	Length of planning horizon	
$s_t \in S$	Patient state consisting of demographics, clinical observations and health state	[148]
$a_t \in A$	Available treatment options	[88]
$b_{min}(s_t)$	Minimum allowable SBP threshold	
$r^{CHD}(s_t)$	Pre-treatment one-period risk of a CHD event	[8]
$r^{Stroke}(s_t)$	Pre-treatment one-period risk of a stroke	[8]
ζ^{CHD}	Subsequent CHD event risk scaling factor	[124]
ζ^{Stroke}	Subsequent stroke risk scaling factor	[25]
$RRR_{a_t}^{CHD}(s_t^\pi)$	Relative risk reduction for CHD	[88]
$RRR_{a_t}^{Stroke}(s_t^\pi)$	Relative risk reduction for stroke	[88]
$r_{a_t}^{CHD}(s_t^\pi)$	Post-treatment one-period risk of a CHD event	
$r_{a_t}^{Stroke}(s_t^\pi)$	Post-treatment one-period risk of a stroke	
$\rho^{CHD}(s_t^\pi)$	Fatality likelihood for CHD events	[95],[9], [66]
$\rho^{Stroke}(s_t^\pi)$	Fatality likelihood for stroke	[95], [9], [66]
$\phi(s_t^\pi)$	Total mortality likelihood	[9]
$P(s_{t+1} \mid s_t, a_t^\pi(s_t))$	State transition probability	
$q(s_t)$	Quality of life weight	[55], [119],[120]
$d_{a_t^\pi}$	Treatment disutility per medication	
λ	Discount factor	
π^*	Optimal treatment strategy	
$V_{T+1}(s_{T+1})$	Terminal condition	
$V_t(s_t)$	Optimal value function	

Table 3.7: QALY model inputs and data sources.

Health State	Quality of Life Weight
Healthy	1
History of CHD but no event this period	0.90
History of stroke but no event this period	0.90
History of CHD and stroke but no event this period	0.81
Survival of a CHD event this period	0.88
Survival of a stroke this period	0.67
Death from non-CVD related cause this period	0
Death from a CHD event this period	0
Death from a stroke this period	0
Dead	0

Table 3.8: Quality of life weights for each health state

Fatality likelihoods for cardiovascular events are presented in Table 3.11.

Health State	**Quality of Life Weight**	**Mortality Scaling Factor**
Healthy	1	1
History of CHD but no event this period	0.90	0.8
History of stroke but no event this period	0.90	0.8
History of CHD and stroke but no event this period	0.81	0.8
Survival of a CHD event this period	0.90	0.8
Survival of a stroke this period	0.90	0.8
Death from non-CVD related cause this period	0	0
Death from a CHD event this period	0	0
Death from a stroke this period	0	0
Dead	0	0

Table 3.9: Quality of life weight and mortality scale factor for terminal health states

Age	**CHD**	**Stroke**
40-49	0.49	0.36
50-59	0.5	0.38
60-69	0.54	0.43
70-79	0.6	0.5
80+	0.67	0.67

Table 3.10: Risk slopes by age and cardiovascular event type

	CHD		**Stroke**	
Age	Male	Female	Male	Female
40-44	0.12	0.26	0.03	0.04
45-54	0.19	0.37	0.07	0.08
55-64	0.26	0.48	0.11	0.11
65-74	0.35	0.64	0.16	0.16
75+	0.44	0.64	0.21	0.21

Table 3.11: Fatality likelihoods by age, sex and cardiovascular event type

3.10.3 Additional Analyses for Time-Varying Benefit and Disutility

Tables 3.12 through 3.20 present the results of the time-varying treatment benefit and disutility assumptions on the second, third and bottom quartiles of age-adjusted CV risk patients. For the treatment disutilities, 0.05^* and 0.1^* denote the increasing disutilities.

Bottom Quartile Age-Adjusted CV Risk Patients	**Rx years per 1000 ppl**				
Treatment Disutility	0.001	0.005	0.01	0.05*	0.1*
Current risk strategy					
No increase in RRR					
5% CV risk	15756	15756	15756	15756	15756
10% CV risk	4721.3	4721.3	4721.3	4721.3	4721.3
15% CV risk	950.07	950.07	950.07	950.07	950.07
20% CV risk	454.56	454.56	454.56	454.56	454.56
10 years till full benefit					
5% CV risk	15794	15794	15794	15794	15794
10% CV risk	4731.3	4731.3	4731.3	4731.3	4731.3
15% CV risk	958.4	958.4	958.4	958.4	958.4
20% CV risk	462.28	462.28	462.28	462.28	462.28
20 years till full benefit					
5% CV risk	15779	15779	15779	15779	15779
10% CV risk	4729.4	4729.4	4729.4	4729.4	4729.4
15% CV risk	957.04	957.04	957.04	957.04	957.04
20% CV risk	460.5	460.5	460.5	460.5	460.5
30 years till full benefit					
5% CV risk	15771	15771	15771	15771	15771
10% CV risk	4727.4	4727.4	4727.4	4727.4	4727.4
15% CV risk	954.95	954.95	954.95	954.95	954.95
20% CV risk	458.4	458.4	458.4	458.4	458.4
Future risk strategy					
No increase in RRR					
5% CV risk	25174	25174	25174	25174	25174
10% CV risk	11084	11084	11084	11084	11084
15% CV risk	2547.2	2547.2	2547.2	2547.2	2547.2
20% CV risk	791.62	791.62	791.62	791.62	791.62
10 years till full benefit					
5% CV risk	25243	25243	25243	25243	25243
10% CV risk	11111	11111	11111	11111	11111
15% CV risk	2558.1	2558.1	2558.1	2558.1	2558.1
20% CV risk	800.64	800.64	800.64	800.64	800.64
20 years till full benefit					
5% CV risk	25228	25228	25228	25228	25228
10% CV risk	11101	11101	11101	11101	11101
15% CV risk	2556.5	2556.5	2556.5	2556.5	2556.5
20% CV risk	800.02	800.02	800.02	800.02	800.02
30 years till full benefit					
5% CV risk	25212	25212	25212	25212	25212
10% CV risk	11097	11097	11097	11097	11097
15% CV risk	2555.2	2555.2	2555.2	2555.2	2555.2
20% CV risk	798.68	798.68	798.68	798.68	798.68

Table 3.12: Treatment (Rx) years for the bottom quartile age-adjusted CV risk patients under different assumptions on time-varying treatment benefit and disutility.

Bottom Quartile Age-Adjusted CV Risk Patients	**CV events per 1000 ppl**				
Treatment Disutility	0.001	0.005	0.01	0.05*	0.1*
Current risk strategy					
No increase in RRR					
5% CV risk	59.11	59.11	59.11	59.11	59.11
10% CV risk	69.26	69.26	69.26	69.26	69.26
15% CV risk	75.29	75.29	75.29	75.29	75.29
20% CV risk	76.57	76.57	76.57	76.57	76.57
10 years till full benefit					
5% CV risk	44.19	44.19	44.19	44.19	44.19
10% CV risk	64.20	64.20	64.20	64.20	64.20
15% CV risk	72.21	72.21	72.21	72.21	72.21
20% CV risk	73.65	73.65	73.65	73.65	73.65
20 years till full benefit					
5% CV risk	48.48	48.48	48.48	48.48	48.48
10% CV risk	65.73	65.73	65.73	65.73	65.73
15% CV risk	72.54	72.54	72.54	72.54	72.54
20% CV risk	74.09	74.09	74.09	74.09	74.09
30 years till full benefit					
5% CV risk	51.91	51.91	51.91	51.91	51.91
10% CV risk	66.48	66.48	66.48	66.48	66.48
15% CV risk	73.15	73.15	73.15	73.15	73.15
20% CV risk	74.79	74.79	74.79	74.79	74.79
Future risk strategy					
No increase in RRR					
5% CV risk	55.51	55.51	55.51	55.51	55.51
10% CV risk	62.57	62.57	62.57	62.57	62.57
15% CV risk	72.57	72.57	72.57	72.57	72.57
20% CV risk	75.75	75.75	75.75	75.75	75.75
10 years till full benefit					
5% CV risk	34.63	34.63	34.63	34.63	34.63
10% CV risk	50.62	50.62	50.62	50.62	50.62
15% CV risk	67.41	67.41	67.41	67.41	67.41
20% CV risk	72.26	72.26	72.26	72.26	72.26
20 years till full benefit					
5% CV risk	37.53	37.53	37.53	37.53	37.53
10% CV risk	54.30	54.30	54.30	54.30	54.30
15% CV risk	68.65	68.65	68.65	68.65	68.65
20% CV risk	72.60	72.60	72.60	72.60	72.60
30 years till full benefit					
5% CV risk	41.41	41.41	41.41	41.41	41.41
10% CV risk	56.67	56.67	56.67	56.67	56.67
15% CV risk	69.26	69.26	69.26	69.26	69.26
20% CV risk	72.97	72.97	72.97	72.97	72.97

Table 3.13: Cardiovascular (CV) disease events for the bottom quartile age-adjusted CV risk patients under different assumptions on time-varying treatment benefit and disutility.

Bottom Quartile Age-Adjusted CV Risk Patients	**QALYs per 1000 ppl**				
Treatment Disutility	0.001	0.005	0.01	0.05*	0.1*
Current risk strategy					
No increase in RRR					
5% CV risk	25687	25665	25637	25654	25631
10% CV risk	25661	25655	25648	25655	25651
15% CV risk	25648	25647	25645	25645	25642
20% CV risk	25646	25645	25645	25644	25642
10 years till full benefit					
5% CV risk	25722	25700	25672	25688	25666
10% CV risk	25670	25664	25657	25664	25660
15% CV risk	25653	25652	25650	25650	25647
20% CV risk	25651	25650	25649	25649	25647
20 years till full benefit					
5% CV risk	25710	25688	25660	25677	25654
10% CV risk	25667	25661	25654	25661	25657
15% CV risk	25652	25651	25650	25649	25647
20% CV risk	25650	25649	25649	25648	25646
30 years till full benefit					
5% CV risk	25702	25680	25652	25669	25646
10% CV risk	25665	25660	25653	25660	25656
15% CV risk	25651	25650	25649	25648	25646
20% CV risk	25649	25648	25647	25647	25645
Future risk strategy					
No increase in RRR					
5% CV risk	25696	25654	25601	25564	25451
10% CV risk	25678	25663	25644	25656	25641
15% CV risk	25653	25650	25646	25646	25640
20% CV risk	25647	25646	25645	25642	25638
10 years till full benefit					
5% CV risk	25751	25709	25656	25619	25505
10% CV risk	25704	25689	25669	25682	25667
15% CV risk	25662	25658	25654	25655	25649
20% CV risk	25653	25652	25651	25648	25643
20 years till full benefit					
5% CV risk	25741	25699	25646	25609	25495
10% CV risk	25695	25679	25660	25673	25658
15% CV risk	25660	25656	25652	25653	25646
20% CV risk	25652	25651	25650	25648	25643
30 years till full benefit					
5% CV risk	25730	25688	25635	25597	25484
10% CV risk	25689	25674	25655	25668	25653
15% CV risk	25659	25655	25651	25651	25645
20% CV risk	25652	25651	25649	25647	25642

Table 3.14: Quality-adjusted life years (QALYs) for the bottom quartile age-adjusted CV risk patients under different assumptions on time-varying treatment benefit and disutility.

3rd Quartile Age-Adjusted CV Risk Patients	**Rx years per 1000 ppl**				
Treatment Disutility	0.001	0.005	0.01	0.05*	0.1*
Current risk strategy					
No increase in RRR					
5% CV risk	24017	24017	24017	24017	24017
10% CV risk	14359	14359	14359	14359	14359
15% CV risk	6565.8	6565.8	6565.8	6565.8	6565.8
20% CV risk	2575.6	2575.6	2575.6	2575.6	2575.6
10 years till full benefit					
5% CV risk	24171	24171	24171	24171	24171
10% CV risk	14439	14439	14439	14439	14439
15% CV risk	6610.8	6610.8	6610.8	6610.8	6610.8
20% CV risk	2616.4	2616.4	2616.4	2616.4	2616.4
20 years till full benefit					
5% CV risk	24132	24132	24132	24132	24132
10% CV risk	14413	14413	14413	14413	14413
15% CV risk	6604.8	6604.8	6604.8	6604.8	6604.8
20% CV risk	2611.9	2611.9	2611.9	2611.9	2611.9
30 years till full benefit					
5% CV risk	24098	24098	24098	24098	24098
10% CV risk	14400	14400	14400	14400	14400
15% CV risk	6596.8	6596.8	6596.8	6596.8	6596.8
20% CV risk	2603	2603	2603	2603	2603
Future risk strategy					
No increase in RRR					
5% CV risk	33823	33823	33823	33823	33823
10% CV risk	23684	23684	23684	23684	23684
15% CV risk	14419	14419	14419	14419	14419
20% CV risk	6494	6494	6494	6494	6494
10 years till full benefit					
5% CV risk	34025	34025	34025	34025	34025
10% CV risk	23835	23835	23835	23835	23835
15% CV risk	14504	14504	14504	14504	14504
20% CV risk	6547	6547	6547	6547	6547
20 years till full benefit					
5% CV risk	34002	34002	34002	34002	34002
10% CV risk	23797	23797	23797	23797	23797
15% CV risk	14480	14480	14480	14480	14480
20% CV risk	6538.7	6538.7	6538.7	6538.7	6538.7
30 years till full benefit					
5% CV risk	33967	33967	33967	33967	33967
10% CV risk	23766	23766	23766	23766	23766
15% CV risk	14468	14468	14468	14468	14468
20% CV risk	6533.6	6533.6	6533.6	6533.6	6533.6

Table 3.15: Treatment (Rx) years for the 3rd quartile age-adjusted CV risk patients under different assumptions on time-varying treatment benefit and disutility.

3rd Quartile Age-Adjusted CV Risk Patients	**CV events per 1000 ppl**				
Treatment Disutility	0.001	0.005	0.01	0.05*	0.1*
Current risk strategy					
No increase in RRR					
5% CV risk	139.96	139.96	139.96	139.96	139.96
10% CV risk	152.02	152.02	152.02	152.02	152.02
15% CV risk	169.08	169.08	169.08	169.08	169.08
20% CV risk	181.28	181.28	181.28	181.28	181.28
10 years till full benefit					
5% CV risk	88.60	88.60	88.60	88.60	88.60
10% CV risk	116.54	116.54	116.54	116.54	116.54
15% CV risk	148.77	148.77	148.77	148.77	148.77
20% CV risk	165.90	165.90	165.90	165.90	165.90
20 years till full benefit					
5% CV risk	96.68	96.68	96.68	96.68	96.68
10% CV risk	126.70	126.70	126.70	126.70	126.70
15% CV risk	153.16	153.16	153.16	153.16	153.16
20% CV risk	167.11	167.11	167.11	167.11	167.11
30 years till full benefit					
5% CV risk	106.81	106.81	106.81	106.81	106.81
10% CV risk	132.95	132.95	132.95	132.95	132.95
15% CV risk	155.65	155.65	155.65	155.65	155.65
20% CV risk	169.31	169.31	169.31	169.31	169.31
Future risk strategy					
No increase in RRR					
5% CV risk	135.87	135.87	135.87	135.87	135.87
10% CV risk	140.23	140.23	140.23	140.23	140.23
15% CV risk	152.09	152.09	152.09	152.09	152.09
20% CV risk	169.11	169.11	169.11	169.11	169.11
10 years till full benefit					
5% CV risk	76.18	76.18	76.18	76.18	76.18
10% CV risk	89.23	89.23	89.23	89.23	89.23
15% CV risk	115.38	115.38	115.38	115.38	115.38
20% CV risk	144.69	144.69	144.69	144.69	144.69
20 years till full benefit					
5% CV risk	80.19	80.19	80.19	80.19	80.19
10% CV risk	97.25	97.25	97.25	97.25	97.25
15% CV risk	124.46	124.46	124.46	124.46	124.46
20% CV risk	150.24	150.24	150.24	150.24	150.24
30 years till full benefit					
5% CV risk	87.05	87.05	87.05	87.05	87.05
10% CV risk	106.99	106.99	106.99	106.99	106.99
15% CV risk	130.79	130.79	130.79	130.79	130.79
20% CV risk	152.88	152.88	152.88	152.88	152.88

Table 3.16: Cardiovascular (CV) disease events for the 3rd quartile age-adjusted CV risk patients under different assumptions on time-varying treatment benefit and disutility.

3rd Quartile Age-Adjusted CV Risk Patients	**QALYs per 1000 ppl**				
Treatment Disutility	0.001	0.005	0.01	0.05*	0.1*
Current risk strategy					
No increase in RRR					
5% CV risk	25186	25147	25098	25074	24981
10% CV risk	25150	25131	25106	25122	25102
15% CV risk	25106	25098	25088	25093	25083
20% CV risk	25081	25077	25073	25071	25063
10 years till full benefit					
5% CV risk	25304	25265	25216	25191	25097
10% CV risk	25220	25200	25175	25191	25171
15% CV risk	25139	25131	25121	25126	25115
20% CV risk	25105	25101	25097	25095	25086
20 years till full benefit					
5% CV risk	25279	25240	25191	25167	25072
10% CV risk	25196	25177	25152	25168	25148
15% CV risk	25131	25123	25113	25118	25107
20% CV risk	25102	25099	25095	25092	25083
30 years till full benefit					
5% CV risk	25254	25215	25166	25141	25047
10% CV risk	25183	25163	25139	25155	25134
15% CV risk	25126	25118	25108	25113	25102
20% CV risk	25098	25095	25091	25088	25079
Future risk strategy					
No increase in RRR					
5% CV risk	25195	25129	25047	24918	24665
10% CV risk	25185	25147	25099	25075	24984
15% CV risk	25148	25128	25103	25113	25086
20% CV risk	25104	25095	25084	25085	25069
10 years till full benefit					
5% CV risk	25344	25278	25195	25064	24810
10% CV risk	25301	25262	25214	25191	25098
15% CV risk	25219	25199	25173	25183	25156
20% CV risk	25144	25135	25125	25125	25109
20 years till full benefit					
5% CV risk	25329	25263	25180	25049	24795
10% CV risk	25277	25238	25190	25166	25074
15% CV risk	25198	25178	25152	25162	25135
20% CV risk	25134	25126	25115	25115	25099
30 years till full benefit					
5% CV risk	25307	25241	25158	25028	24774
10% CV risk	25253	25214	25166	25142	25050
15% CV risk	25185	25165	25140	25149	25122
20% CV risk	25130	25121	25110	25110	25094

Table 3.17: Quality-adjusted life years (QALYs) for the 3rd quartile age-adjusted CV risk patients under different assumptions on time-varying treatment benefit and disutility.

2nd Quartile Age-Adjusted CV Risk Patients	**Rx years per 1000 ppl**				
Treatment Disutility	0.001	0.005	0.01	0.05*	0.1*
Current risk strategy					
No increase in RRR					
5% CV risk	28728	28728	28728	28728	28728
10% CV risk	20593	20593	20593	20593	20593
15% CV risk	14088	14088	14088	14088	14088
20% CV risk	8645.5	8645.5	8645.5	8645.5	8645.5
10 years till full benefit					
5% CV risk	29085	29085	29085	29085	29085
10% CV risk	20860	20860	20860	20860	20860
15% CV risk	14273	14273	14273	14273	14273
20% CV risk	8783.8	8783.8	8783.8	8783.8	8783.8
20 years till full benefit					
5% CV risk	29026	29026	29026	29026	29026
10% CV risk	20795	20795	20795	20795	20795
15% CV risk	14230	14230	14230	14230	14230
20% CV risk	8766.5	8766.5	8766.5	8766.5	8766.5
30 years till full benefit					
5% CV risk	28959	28959	28959	28959	28959
10% CV risk	20746	20746	20746	20746	20746
15% CV risk	14203	14203	14203	14203	14203
20% CV risk	8746.1	8746.1	8746.1	8746.1	8746.1
Future risk strategy					
No increase in RRR					
5% CV risk	38562	38562	38562	38562	38562
10% CV risk	30218	30218	30218	30218	30218
15% CV risk	23193	23193	23193	23193	23193
20% CV risk	16759	16759	16759	16759	16759
10 years till full benefit					
5% CV risk	38974	38974	38974	38974	38974
10% CV risk	30588	30588	30588	30588	30588
15% CV risk	23491	23491	23491	23491	23491
20% CV risk	16977	16977	16977	16977	16977
20 years till full benefit					
5% CV risk	38947	38947	38947	38947	38947
10% CV risk	30534	30534	30534	30534	30534
15% CV risk	23426	23426	23426	23426	23426
20% CV risk	16928	16928	16928	16928	16928
30 years till full benefit					
5% CV risk	38895	38895	38895	38895	38895
10% CV risk	30468	30468	30468	30468	30468
15% CV risk	23371	23371	23371	23371	23371
20% CV risk	16897	16897	16897	16897	16897

Table 3.18: Treatment (Rx) years for the 2nd quartile age-adjusted CV risk patients under different assumptions on time-varying treatment benefit and disutility.

2nd Quartile Age-Adjusted CV Risk Patients	**CV events per 1000 ppl**				
Treatment Disutility	0.001	0.005	0.01	0.05*	0.1*
Current risk strategy					
No increase in RRR					
5% CV risk	287.54	287.54	287.54	287.54	287.54
10% CV risk	300.62	300.62	300.62	300.62	300.62
15% CV risk	320.20	320.20	320.20	320.20	320.20
20% CV risk	343.30	343.30	343.30	343.30	343.30
10 years till full benefit					
5% CV risk	166.61	166.61	166.61	166.61	166.61
10% CV risk	198.61	198.61	198.61	198.61	198.61
15% CV risk	239.78	239.78	239.78	239.78	239.78
20% CV risk	282.63	282.63	282.63	282.63	282.63
20 years till full benefit					
5% CV risk	178.84	178.84	178.84	178.84	178.84
10% CV risk	216.11	216.11	216.11	216.11	216.11
15% CV risk	257.94	257.94	257.94	257.94	257.94
20% CV risk	293.89	293.89	293.89	293.89	293.89
30 years till full benefit					
5% CV risk	195.96	195.96	195.96	195.96	195.96
10% CV risk	234.58	234.58	234.58	234.58	234.58
15% CV risk	269.69	269.69	269.69	269.69	269.69
20% CV risk	300.39	300.39	300.39	300.39	300.39
Future risk strategy					
No increase in RRR					
5% CV risk	282.77	282.77	282.77	282.77	282.77
10% CV risk	286.30	286.30	286.30	286.30	286.30
15% CV risk	295.22	295.22	295.22	295.22	295.22
20% CV risk	311.04	311.04	311.04	311.04	311.04
10 years till full benefit					
5% CV risk	151.04	151.04	151.04	151.04	151.04
10% CV risk	162.98	162.98	162.98	162.98	162.98
15% CV risk	185.98	185.98	185.98	185.98	185.98
20% CV risk	220.13	220.13	220.13	220.13	220.13
20 years till full benefit					
5% CV risk	156.35	156.35	156.35	156.35	156.35
10% CV risk	174.05	174.05	174.05	174.05	174.05
15% CV risk	201.71	201.71	201.71	201.71	201.71
20% CV risk	237.51	237.51	237.51	237.51	237.51
30 years till full benefit					
5% CV risk	166.83	166.83	166.83	166.83	166.83
10% CV risk	190.21	190.21	190.21	190.21	190.21
15% CV risk	220.17	220.17	220.17	220.17	220.17
20% CV risk	252.02	252.02	252.02	252.02	252.02

Table 3.19: Cardiovascular (CV) disease events for the 2nd quartile age-adjusted CV risk patients under different assumptions on time-varying treatment benefit and disutility.

2nd Quartile Age-Adjusted CV Risk Patients	**QALYs per 1000 ppl**				
Treatment Disutility	0.001	0.005	0.01	0.05*	0.1*
Current risk strategy					
No increase in RRR					
5% CV risk	24595	24543	24477	24403	24233
10% CV risk	24559	24527	24487	24477	24410
15% CV risk	24507	24487	24463	24469	24439
20% CV risk	24452	24441	24427	24427	24405
10 years till full benefit					
5% CV risk	24844	24791	24726	24649	24477
10% CV risk	24751	24719	24678	24668	24600
15% CV risk	24643	24623	24599	24604	24573
20% CV risk	24544	24533	24519	24518	24494
20 years till full benefit					
5% CV risk	24808	24755	24690	24613	24440
10% CV risk	24709	24677	24636	24625	24557
15% CV risk	24606	24586	24561	24567	24535
20% CV risk	24524	24513	24499	24498	24474
30 years till full benefit					
5% CV risk	24765	24713	24647	24571	24398
10% CV risk	24670	24637	24597	24586	24518
15% CV risk	24583	24564	24539	24544	24513
20% CV risk	24512	24501	24487	24486	24462
Future risk strategy					
No increase in RRR					
5% CV risk	24603	24519	24413	24211	23849
10% CV risk	24598	24542	24471	24380	24185
15% CV risk	24573	24535	24487	24459	24362
20% CV risk	24527	24502	24471	24469	24422
10 years till full benefit					
5% CV risk	24890	24806	24700	24493	24125
10% CV risk	24856	24799	24728	24635	24437
15% CV risk	24785	24747	24699	24669	24571
20% CV risk	24686	24661	24630	24627	24580
20 years till full benefit					
5% CV risk	24872	24787	24682	24475	24107
10% CV risk	24823	24766	24695	24601	24403
15% CV risk	24744	24706	24658	24628	24530
20% CV risk	24648	24623	24592	24589	24541
30 years till full benefit					
5% CV risk	24840	24756	24650	24444	24077
10% CV risk	24781	24724	24653	24560	24362
15% CV risk	24704	24665	24618	24588	24489
20% CV risk	24621	24596	24564	24561	24513

Table 3.20: Quality-adjusted life years (QALYs) for the 2nd quartile age-adjusted CV risk patients under different assumptions on time-varying treatment benefit and disutility.

CHAPTER IV

Copayment Restructuring for Improved Adherence to Treatment Plans

This chapter builds upon the research completed in Chapter III to understand how to incentivize adherence to optimal treatment policies. This work develops a bilevel optimization model based upon a Stackelberg game where the leader is a single insurance provider who determines medication copayments for its insured patients. The followers are patient classes who react to the insurance provider's copayment level by altering their adherence to medication. As discovered in Chapter III, the optimal treatment policy for a patient depends on his/her adherence to medication; therefore, the new copayment level also potentially alters the optimal treatment policy. To solve this nonlinear bilevel model, we propose an iterative scheme which repeatedly solves the lower level problem (determining the optimal treatment policy) which provides inputs to the upper level problem (determining the optimal medication copayments). We applied this model to patients, aged 65 and older found in the NHANES dataset, who may be representative of those patients enrolled in Medicare and compared the health and economic performance of the optimal copayments against current practice.

4.1 Background

Operations research has contributed heavily to the derivation of optimal treatment guidelines for chronic disease [2], [4], [5], [31], [133], [81], [13], [139], [141]. However, this research has assumed that physicians and patients will follow the optimal guidelines, which may not always be the case. As shown in Chapter III, decreased adherence to treatment plansmay negatively impact outcomes. Hence, it is crucial to determine how improved patient adherence may be incentivized.

One source for these incentives is the insurance provider. An insurance provider sets the price a patient pays for medical services, including pharmacotherapy (e.g. hypertension medication). Theoretical and empirical research in health economics [80], [30], [24], [34], [97] suggests that as the price of medication decreases, adherence to the medication will increase due to price elasticity. Therefore, insurance providers may be able to directly address low patient adherence to treatment policies by decreasing copayments (i.e. the price paid per medication). With the principle of price elasticity as the motivation for a financial incentive to improve adherence, we propose a mathematical program for determining optimal copayments paid by patients to their insurance provider.

We consider the perspective of a single insurance provider whose health care plan covers a patient population which is heterogeneous in its risk factors for cardiovascular disease. The insurance provider wishes to estimate what copayment to assign to each patient class (i.e. a cluster of patients similar in their risk factors) while taking into consider that patient heterogeneity leads to differences in optimal treatment plans and variability in the marginal benefit of improved adherence. We assume the insurance provider is a centralized decision maker who is interested in setting

copayment levels for its patients that would maximize the expected total health of its patient population. Free medications for all patients would naturally maximize total health; however, free medications would be very expensive to the insurance provider who is operating in an environment with increasing national expenditures on health care [26] and a strong focus on cost control [67], [114]. On the other hand, the provider may wish to target those patients who would benefit the most. We also assume the insurance provider is concerned with the inequity caused by different copayments for different patients classes.

4.2 Literature Review

This research intersects the fields of payment structures, health economics and optimization. Within the domain of operations research and payment structures, there have been studies on outcome-based reimbursement policies, including how threshold reimbursements affect prescription policies for a clinic [147] and the use of principal-agent models for identifying measures for determining reimbursement/incentives between a healthcare provider (e.g. hospital) and the payer (e.g. insurance provider) [89], [56]. Research has focused primarily on contracts between payers and providers: threshold performance-based contract in order to coordinate payer-provider system [72], and coordination of relationships between payer-provider (maximize both payer's and provider's objectives) in preventive health care through ranking patients with a threshold for providing treatment [157]. This research has also extended into copayment coupons from drug manufacturers and the desirability of copay coupon bans to insurers [75]. In addition to the study of pricing between hospitals and insurance providers, researchers have studied pricing mechanisms between manufacturers and insurance providers [162], [160], [149], [96]. Our research differs in that

we model the relationship between patients (the end-user of services) and the insurance provider (the primary payer for services). Furthermore, our research considers treatment plans (prescribed by the healthcare provider) which are sensitive to the patient's characteristics and adherence to medication. The focus on the patient requires the modeling of disease dynamics over time and an optimization model for deriving treatment plans. Our research also incorporates inequity of copayment coverage, which differs from current research that seeks a purely utilitarian solution.

The field of health economics has also studied copayments, including stylistic models for determining whether a new health technology should be subsidized [145] or whether medical interventions should be included in statutory insurance packages [146]. Health economists have considered the welfare economics of public drug insurance and government subsidies for health care premiums [86]. They have also researched the impact of payments and altruistic physicians on treatment decisions between physician and patient [94]. Other research looks at financial incentives to control treatment decisions [29], which considers inexpensive vs. expensive treatments for an observable disease state, as well as economic models for copayment reduction [35], [33], [129]. More empirical research includes studying the impact of adherence on health [125], [104] as well as real-world implementations of cost-sharing changes [97], [30], [24], [34]. While the previous research has addressed the economics of copayments, our work differs from this research in that we employ optimization to determine copayment levels which are tailored to the patient's characteristics and his/her optimal treatment plan.

To determine the optimal copayment levels, we formulate the problem as a bilevel optimization model (see 4.3.1 for details) with a single leader (i.e. insurance provider) and multiple followers (i.e. patient classes). Hierarchical models, including bilevel,

are generally challenging to solve due to the interplay between the leader and followers [158]. Traditional approaches to solving hierarchical models include the Kuhn Tucker approach [62], [43] and penalty functions [153]. More recent solution techniques include genetic algorithms [92], evolutionary algorithms [144], and tabu search [165]. Researchers have found that there is no single best approach to solving these models, and that heuristic approaches combined with problem domain expertise are more robust [115] than the traditional methods. We continue in this trend of heuristic algorithms and develop an iterative nonlinear programming approach which utilizes the findings of Chapter III to solve the followers' optimization problem.

Overall, our research has the following key contributions: (1) a bilevel optimization formulation for determining optimal copayments for different patients, (2) the direct modeling of inequity constraints, (3) an iterative algorithm for solving the bilevel problem, and (4) a comparative analysis between optimal copayments and the current Medicare system. To our knowledge, this is the first study to integrate optimal copayments and optimal treatment plans for hypertension. While this work is applied to hypertension, the formulation and modeling principles are applicable to other chronic diseases for which patients are required to pay out-of-pocket to treat their disease.

4.3 Model

Figure 4.1 presents the dynamics and relationships between decisions and outcomes in the insurance/patient healthcare system. Our decision maker is a single insurance provider servicing a heterogeneous patient population. We assume the population of patients can be segmented into different classes. These classes may be defined based on behavior, e.g. smoking vs. nonsmoking, and/or cardiovascular

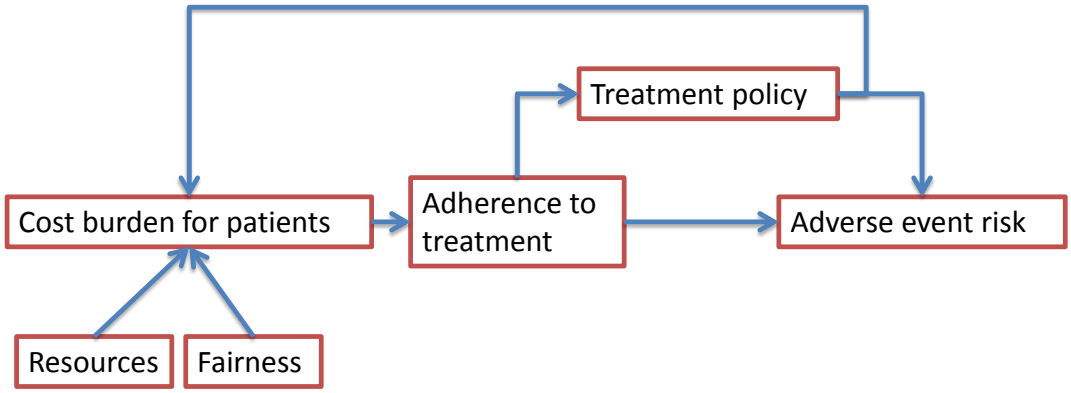

Figure 4.1: Flow diagram capturing the dynamics of the health system and the relationship between coverage decisions and outcomes.

risk, e.g. the patient's diabetes status. Let $p \in P$ denote a particular patient class and let n_{pht} be the expected number of patients in class p and health state $h \in H$ (as defined in 3.8) at each time period $t \in T$. The benevolent insurance provider is interested in determining the medication coverage $y_{ph} \in [0, 1]$, defined as the percentage of medication cost covered by the insurance provider, for each patient class p that maximizes total population health. We assume the insurance provider does not want the copayment level for a particular patient class to change from year to year within the planning horizon, since this would create implementation challenges with its patients. Assuming a quality of life $q_{pht} \in [0, 1]$ and discount factor $\lambda \in [0, 1]$, the insurance provider's objective function is of the form:

$$Z^* = \max_{\mathbf{y}} \left\{ \sum_{t=1}^{T} \sum_{p \in P} \sum_{h \in H} \lambda^t \cdot n_{pht}(\mathbf{y}) \cdot q_{pht} \right\} \tag{4.1}$$

where $\mathbf{y}$ is the matrix of coverage decisions and Z^* is the optimal objective function value.

Using the theory of price elasticity, we assume the following relationship between

the coverage decision y_{ph} and the patient's adherence to medication α_{pht}:

$$\alpha_{pht} = \max\{0, \min\{1, \alpha^i_{pht} + \delta^\alpha_{pht} \cdot (y_{ph} - y^i_{ph})\}\} \tag{4.2}$$

where δ^α_{pht} is the price elasticity of adherence, i.e. the increase in adherence caused by a 1% reduction in the price of medication; α^i_{pht} is the current adherence rate and y^i_{ph} is the current coverage. When δ^α_{pht} is positive and the change in coverage is positive (i.e. new coverage is higher than current coverage), the patient's adherence to medication will increase.

From our analyses in Chapter III, we found that the optimal treatment policy depends on the patient's particular cardiovascular risk factors and adherence to medication. Let π^*_{pht} be the optimal treatment for patient class p in health state h at time t as determined by the MDP (see equation (3.19)). In addition to providing the optimal treatment policy, the MDP also implicitly computes the probability that a patient of class $p \in P$ will be in each of the health states $h \in H$ at each time period $t \in T$. Using the state transition probabilities, we can compute the expected number of patients in each combination of health state and time for a particular patient class, n_{pht}.

Given the drug manufacturer's price $c(\cdot)$, the cost of the optimal treatment to the patient is $(1 - y_{ph}) \cdot c(\pi^*_{pht})$ and the cost of the optimal treatment to the insurance provider is $y_{ph} \cdot c(\pi^*_{pht})$. The insurance provider is assumed to have a limited budget each time period, B_t, defined on a per patient basis, which constrains the total cost-sharing relief available from the insurance provider. We assume this budget represents the additional amount of resources the insurance provider is willing to expend to reduce copayments for its patients. We do not consider extracting resources from other segments of the insurance provider's budget to pay for copayment reductions.

We can write the insurance provider's resource constraint as:

$$\sum_{p\in P}\sum_{h\in H} n_{pht}(\mathbf{y})\cdot y_{ph}\cdot c(\pi^*_{pht}) \leq B_t \cdot \sum_{p\in P}\sum_{h\in H} n_{pht}(\mathbf{y}) \tag{4.3}$$

We further assume that the benevolent insurance provider wants to limit the total inequity of copayment coverage present in the allocation of resources across the different patient classes. To compute the total allocation inequity, we propose to use a social welfare function $W(\mathbf{y})$ paired with a maximum inequity threshold $w \in [0, 1]$. We will use the Theil index as our social welfare function $W(\mathbf{y})$. The Theil index measures the degree of entropy or disorder in a system, and it is a special case of the generalized entropy ratio [110]. A low Theil index corresponds to an equitable distribution of resources (i.e. low disorder) and a high Theil index corresponds to an inequitable distribution (i.e. high disorder). We compute the Theil index as:

$$W(\mathbf{y}) = \frac{1}{N\ln N}\sum_{p\in P}\sum_{h\in H}\frac{y_{ph}}{\bar{y}}\ln\frac{y_{ph}}{\bar{y}} \leq w \tag{4.4}$$

$$\bar{y} = \frac{1}{N}\sum_{p\in P}\sum_{h\in H} y_{ph} \tag{4.5}$$

$$N = |P|\cdot|H| \tag{4.6}$$

where N is the total number of combinations of patient class and health state and $\bar{y}$ is the average coverage across all patient classes. The scale factor $\frac{1}{N\ln N}$ ensures that the Theil index takes values between 0 and 1. Note that a patient class with no difference from average coverage (i.e $y_{ph} = \bar{y}$) does not contribute to the index since $1 \cdot \ln(1) = 0$. Therefore, when all copayment coverages are equal, the Theil index equals 0. As coverages begin to deviate from the average coverage, the Theil index increases toward 1. Therefore, with a maximum inequity threshold w, we can appropriately limit the total allocation inequity to match the insurance provider's desired level of fairness.

Formally, we can write the mathematical program as follows:

$$Z^* = \max_{\mathbf{y}} \left\{ \sum_{t=1}^{T} \sum_{p \in P} \sum_{h \in H} \lambda^t \cdot n_{pht}(\mathbf{y}) \cdot q_{pht} \right\} \tag{4.7}$$

$$\text{s.t.} \quad \sum_{p \in P} \sum_{h \in H} n_{pht}(\mathbf{y}) \cdot y_{ph} \cdot c_(\pi^*_{pht}) \leq B_t \cdot \sum_{p \in P} \sum_{h \in H} n_{pht}(\mathbf{y}) \tag{4.8}$$

$$W(\mathbf{y}) = \frac{1}{N \ln N} \sum_{p \in P} \sum_{h \in H} \frac{y_{ph}}{\bar{y}} \ln \frac{y_{ph}}{\bar{y}} \leq w \tag{4.9}$$

$$\alpha_{pht} = \max\{0, \min\{1, \alpha^i_{pht} + \delta^{\alpha}_{pht} \cdot (y_{ph} - y^i_{ph})\}\} \tag{4.10}$$

$$y_{ph} \in [0, 1] \tag{4.11}$$

From the formulation above, we see that the optimal treatment policy π^*_{pht} is embedded within the constraint (4.8). Furthermore, the optimal treatment policy depends on the patient's adherence which is a function of the coverage decision variable y_{ph}. Additionally, the expected number of patients n_{pht}, as computed by the MDP under the optimal treatment policy, is present in both the objective function and the resource constraint. This mathematical program shows the complex nature of the insurance provider/patient relationship and the necessary hierarchy of the optimization model.

4.3.1 Bilevel Optimization

The hierarchical structure of the copayment problem can be characterized as a Stackelberg game [16], [163]. A Stackelberg game involves a leader and follower, where the leader makes a decision and the follower reacts to that decision. The leader is assumed to have knowledge of the follower's reaction when making his decision. When both parties act optimally, the problem can be formulated as a bilevel optimization problem (BLP). The general formulation for a BLP is:

$$\max_{\mathbf{y},\boldsymbol{\pi}(\mathbf{y})} \quad F(\mathbf{y}, \boldsymbol{\pi}(\mathbf{y})) \tag{4.12}$$

$$\text{s.t.} \quad G_i(\mathbf{y}, \boldsymbol{\pi}(\mathbf{y})) \leq 0 \quad \forall i \tag{4.13}$$

$$\pi_p(\mathbf{y}) = \arg\max_{\pi} f_p(\mathbf{y}, \pi) \quad \forall p \tag{4.14}$$

$$g_{pj}(\mathbf{y}, \pi_p) \leq 0 \quad \forall p, j \tag{4.15}$$

The general formulation applies to the mathematical program described by equations (4.7) to (4.11). In the BLP, F corresponds to (4.7) while the set of G_i correspond to the constraints on resources (4.8) and inequity (4.9). Similarly, π_p matches the optimal treatment policy π^*_{pht} for each patient class p and g_{pj} denotes the SBP constraint for the selected treatments as defined in Section 3.3. As mentioned in 4.2, BLPs are difficult to solve due to the complex relationship between leader and followers. To solve the BLP, we propose an iterative nonlinear programming technique.

4.4 Solution Methodology

Our algorithm iteratively solves the follower's optimization problems (modeled as the MDP from Chapter III) to provide information to the leader. The leader then uses this information to determine copayment levels. The new copayment levels and corresponding adherence rates are passed back to the followers and the MDPs are resolved. The algorithm repeats this cycle of information flow between followers and the leader until the objective function Z^* has not sufficiently changed. The specifics of the algorithm are described below:

1. INITIALIZE: Choose convergence parameter ϵ. Set $i = 1$, $Z_i = \infty$, $\mathbf{y}_0 = \mathbf{y}^i$ and $\alpha_{pht} = \alpha^i_{pht}$.

2. Solve the MDP for all patient classes $p \in P$ under α_{pht} to obtain π^*_{pht}. Set $\pi_{pht} = \pi^*_{pht}$.

3. Solve the NLP under π_{pht} to obtain $\mathbf{y}^*$, α^*_{pht} and Z^*. Store $\mathbf{y}_i = \mathbf{y}^*$ and set $\alpha_{pht} = \alpha^*_{pht}$ and $Z_{i+1} = Z^*$.

4. IF $Z_{i+1} \in [(1-\epsilon) \cdot Z_i, (1+\epsilon) \cdot Z_i]$ OR $\mathbf{y}_i = \mathbf{y}_j$ for some $j < i$ THEN Stop ELSE Set $i = i + 1$ and return to Step 2 .

The termination condition for the iterative algorithm requires that the convergence condition be satisfied, i.e. the objective function value is within $\pm\epsilon$ of the previous iterate's objective function value, or if the copayment coverage matrix $\mathbf{y}$ has been revisited. The second termination condition is an anti-cycling rule akin to the anti-cycling rules used in the simplex algorithm for linear optimization. Since the set of optimal treatment policies is finite $\pi^* \in \Pi, |\Pi| < \infty$, there are finitely many iterations of algorithm before the anti-cycling rule applies or the convergence condition is satisfied.

4.5 Results

To numerically test our model and algorithm, we considered the perspective of Medicare, the public insurance provider for senior citizens in the United States. We assumed Medicare is interested in optimizing their copayments for hypertension medication, and that all other copayments and deductibles remained constant. Table 4.1 summarizes the key inputs to the BLP.

We parameterized the model using assuming a population of 10,000 senior citizens (those aged 65 and older) who were randomly sampled from NHANES III, survey data comprising a representative sample of the U.S. population [148]. We believe this sample population is a good representation of patients who are enrolled

Input	Value	Source
Number of patients	10,000	
Patient data	NHANES III	[148]
Current coverage	37%	[35]
Current adherence	50%	[150]
Price elasticity of adherence	0.16	[37]
Risk calculator	Framingham	[8]
Treatment benefit	Meta-analysis of RCTs	[88]
Mortality and fatality likelihoods	CV event and other-cause	[78], [9]
CV event costs	$12,229	[35]
Drug cost per year	$212	[35]

Table 4.1: BLP inputs and data sources

in Medicare. We assumed a 10 year planning horizon. In our analysis, we considered four patient classes: diabetic smokers, diabetic nonsmokers, nondiabetic smokers, and nondiabetic nonsmokers with the following prevalences: 2.3%, 23.6%, 6.7%, and 67.4%, respectively [52], [11]. Currently under Medicare, 37% of the cost of hypertension medication is covered by standard insurance. We assumed all patients had a current adherence α^i of 50%. Based on health economics literature, we assumed a price elasticity of 0.16 which indicates a relatively low change in adherence due to copayment changes [37]. All sources for cardiovascular risk calculations and mortality and fatality likelihoods are from the analysis in 3.8. Hypertension drug costs and cardiovascular event costs were computed as average costs.

Figure 4.2 shows the BLP coverage policy at different budget levels when there is no constraint on inequity, i.e. $w = 1$. The BLP policy prioritizes diabetic patients, both smokers and nonsmokers, over nondiabetic patients in the allocation of coverage resources. Specifically, when Medicare is budget-neutral ($B_t = 0$), the BLP policy provides free hypertension medication to diabetic patients. To offset the cost of giving free medication to diabetic patients, Medicare must increase the copayments for nondiabetic patients: 17% coverage for nondiabetic nonsmokers and 0% for nondiabetic smokers. As the budget increases, the BLP policy allocates resources

to the nondiabetic nonsmokers until they receive free medication. Finally, remaining resources are directed to nondiabetic smokers until all patients are at full copayment coverage.

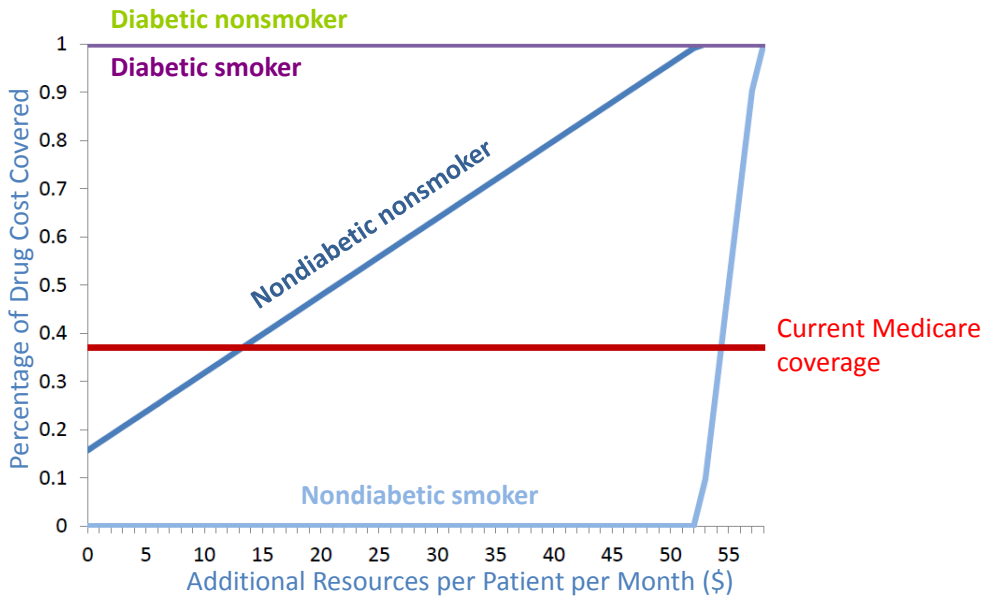

Figure 4.2: Optimal coverage decisions at different budget levels assuming no constraint on inequity.

For the BLP policies described above, we computed the change in QALYs per 1000 patients as the expected QALYs under the BLP policy minus the expected QALYs under the current Medicare coverage. Figure 4.3 graphs the QALY change as a function of the budget. When Medicare is budget-neutral, the change in QALYs per 1000 patients is 4.35 over the next 10 years. At full coverage for all patients, the QALY gain is 15.37 per 1000 patients over the next 10 years.

Figure 4.4 shows the change in cost under the BLP policy, defined as the expected total medication and cardiovascular disease event costs for the BLP policy minus the expected total costs for the current Medicare coverage. At low resource levels, the change in cost is negative (-$36,184.40 per 1000 patients when budget-neutral)

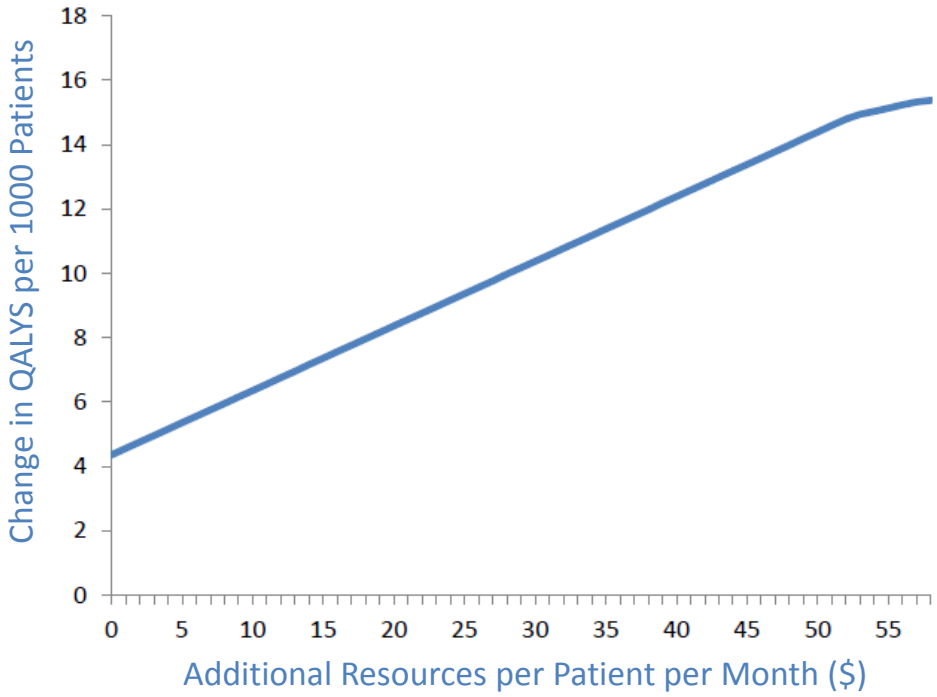

Figure 4.3: Change in quality-adjusted life years (QALYs) with respect to current Medicare coverage.

indicating that the BLP policy is cost-saving. As the resource expenditures increase, the costs of covering medications exceeds the savings from fewer CV events which leads to higher costs under the BLP policy.

The above analyses assumed there was no constraint on the total inequity caused by the BLP allocation. Figure 4.5 shows the performance of the BLP policy in terms of QALY gains as a function of the maximum allowable inequity threshold w, assuming Medicare is budget-neutral. When $w = 0$ (no inequity allowed) and $B_t = 0$ (budget-neutral), the BLP policy is the same as the current Medicare coverage since all patients must receive the same coverage and Medicare has no additional resources to expend for reducing copayments. Therefore, the QALY gain is 0. As

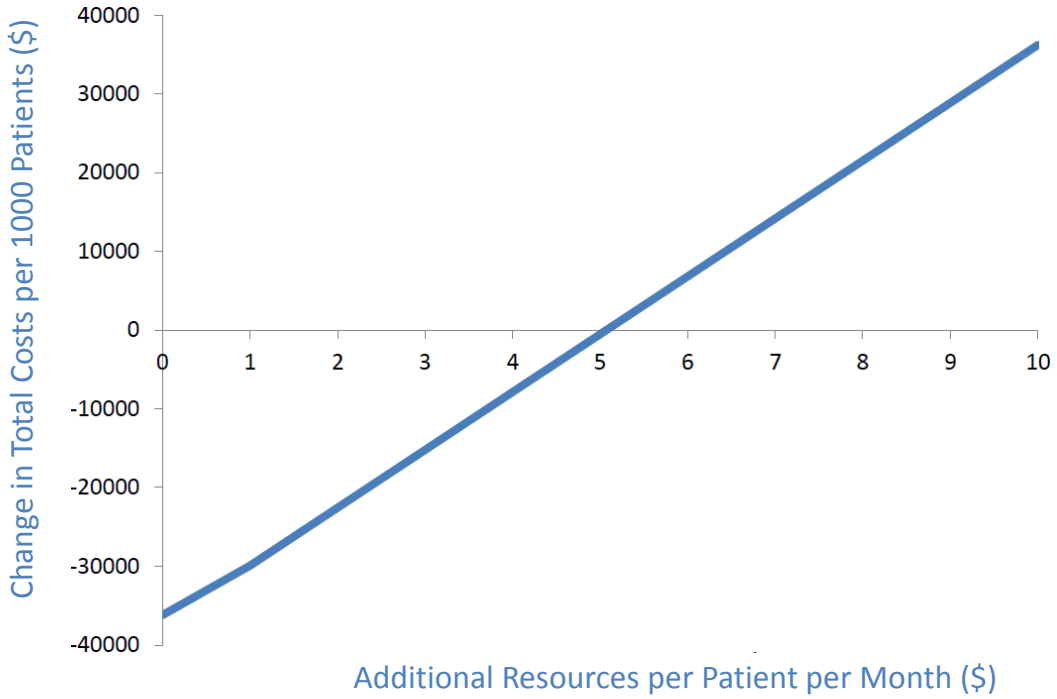

Figure 4.4: Change in cost of optimal coverage with respect to current Medicare coverage.

the maximum allowable inequity is increased, the QALY gain increases toward its maximum (assuming budget-neutral) of 4.35 QALYs per 1000 patients over the next 10 years.

Combining the change in QALYs with the change in costs, we computed an incremental cost-effectiveness ratio (ICER). ICER is defined as the ratio of the change in cost to the change in QALYs for the BLP policy. Hence, a small change in cost and large change in QALYs will yield a low ICER. The lower the ICER for the BLP policy, the more cost-effective the policy becomes. Figure 4.6 presents a heat map of the ICER for the BLP policies under combinations of resource constraints and maximum allowable inequity. Dark colors correspond to low ICERs (more cost-effective) and light colors correspond to high ICERs (less cost-effective). As illustrated, when the maximum allowable inequity decreases, the ICER for the BLP policy increases. Sim-

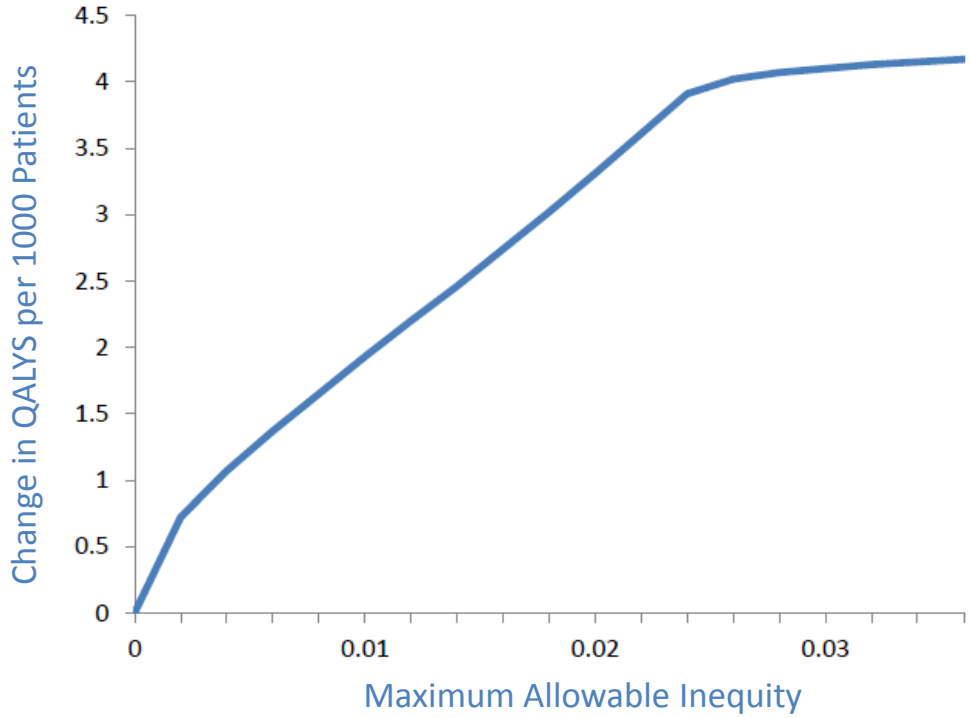

Figure 4.5: QALYs saved under different levels of maximum allowable inequity.

ilarly, as the resource level increases, the ICER increases. Given a cost-effectiveness threshold, Medicare could use Figure 4.6 to determine if the BLP policy for a particular resource level and inequity constraint would be cost-effective.

4.6 Discussion

We have developed a bilevel optimization model to determine copayment levels for individual patient classes within a heterogeneous patient population serviced by a single insurance provider. In this model, the leader is the insurance provider whose objective is to maximize total population health by choosing a percentage of the medication cost to cover for each patient class. The patients (followers) then react to the leader's coverage decision by adjusting their adherence to medication and

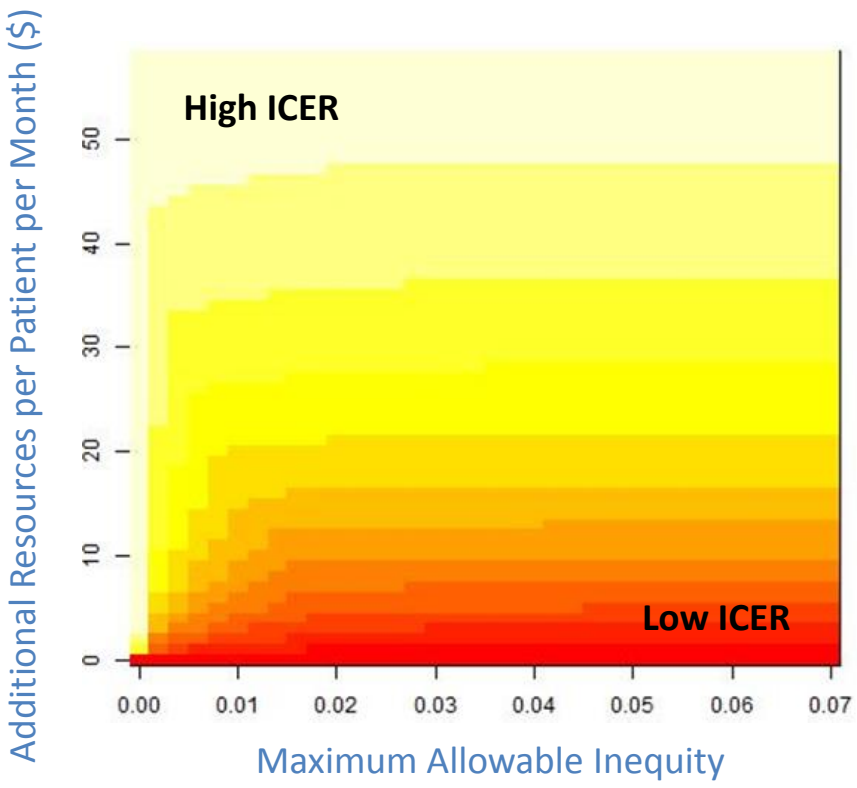

Figure 4.6: Incremental cost-effectiveness ratio (ICER) at different budgets and inequity levels.

their optimal treatment policy, which depends on their adherence rate and disease characteristics. We developed an iterative nonlinear programming algorithm which exploits the leader/follower relationship to derive locally optimal solutions to the BLP.

In our case-study, we found that the BLP policy prioritizes diabetic patients over nondiabetic patients. Diabetic patients are at higher risk for cardiovascular disease events and thus receive greater benefit from improved adherence as a result of lower copayments. Amongst nondiabetic patients, the BLP policy prioritizes nonsmokers over smokers because smokers are at higher risk for non-CV related death, such as lung cancer. Hence, nondiabetic smokers have the smallest marginal benefit of

improved adherence.

When the insurance provider is able to seek the most efficient allocation of resources, i.e. he has no consideration of inequity when deciding coverages, we found that inequity is maximized at low resource budgets. The BLP policy prioritizes diabetic patients which requires the recuperation of costs via increased copayments for nondiabetic patients. However, as the resource budget increases, ceiling effects and diminishing returns on those resources lead to lower total inequity in the copayment coverages. We also found that fairness can be relatively costly, but restrictions on the total inequity caused by the coverage may be a necessary consideration for a public institution like Medicare. However, when the insurance provider has more resources available for expenditure the cost of fairness is lower.

The health benefits from the BLP policy depend strictly on the ability of copayment restructuring to improve adherence to medication for high risk patients. The relatively small improvement in health (15.37 QALYs per 1000 patients) in our numerical study is a reflection of the small price elasticity of adherence δ^{α} found in the health economics literature. The change in coverage from 37% to 100% leads to a 10% increase in adherence to medication under the assumed price elasticity. This small improvement in adherence generates a small change in QALYs for a single patient, as illustrated in Chapter III. If certain patients were more sensitive to the price of medications, the health gains would increase accordingly.

Overall, we found that the cost-effectivess, as defined by ICER, is decreasing with the resource budget and decreasing as the maximum allowable inequity threshold decreases. As more resources are expended, the change in total costs is increasing, but at a faster rate than the change in QALYs due to diminishing returns. Therefore, ICER increases as the budget increases. And we found that at low resource budgets,

the BLP policy is cost-saving compared to current Medicare coverage. Similarly, as the maximum allowable inequity decreases, the BLP policy is forced to allocate resources to low priority patients which could go to high priority patients. By enforcing more equal allocations across the population, efficiency is lost but fairness is gained. This loss of efficiency leads to less cost-effective BLP policy.

4.6.1 Limitations

One limitation of this work is that our proposed iterative nonlinear programming algorithm does not guarantee a globally optimal solution. Similar to other research in bilevel optimization, heuristic approaches have been used to solve the models. We may consider other heuristic algorithms, such as genetic algorithms, but these do not resolve the issue of global optimality. Another approach may involve linear approximations of the BLP presented in order to obtain globally optimal solutions to the approximate problem. With uncertainty around the modeling inputs, optimally solving the approximate problem may be better than local optimal solutions to the exact problem.

In our numerical analysis, we considered as a case-study only a small number of patient classes: combinations of smoking and diabetes status. The full Medicare population is comprised of a richer set of classes including differences in socioeconomic status, sex, and race. These additional features of the patients alter both the optimal treatment policy and their price elasticity of adherence. Future research remains in extending the numerical study to include more patient classes to better represent the heterogeneity present within the Medicare population.

We restricted the insurance provider's mechanism for improving adherence to fiscal incentives in the form of medication copayment reduction. Numerically, we found that the small price elasticity of adherence leads to small health improvements.

This suggests other mechanisms may be useful for improving adherence to optimal treatment policies, such as implementing educational programs about the value of medication and registering patients with a text-message system to remind them to take their medication as prescribed. Our modeling framework is amenable to extending the decision space to include a portfolio of incentives, where the insurance provider can select the intensity of the incentive (e.g. the number of text messages to send) for each patient class. Using a similar modeling strategy as the financial incentive, the impact of the portfolio of incentives on the patient's adherence can be structured using correlated adherence elasticities for each incentive.

We also assumed our insurance provider is benevolent, i.e. the insurance provider's objective was to maximize the total health of its population. This assumption holds true for a public insurance provider like Medicare. However, a private insurance provider may be more interested in minimizing cost. We believe that our inclusion of a resource constraint in the BLP captures the key elements important to a cost-minimizing decision maker, but directly altering the objective function to minimize total expected costs is an intriguing research extension for this work.

Lastly, we assumed perfect information on the part of the insurance provider. In practice, the insurance provider may not know exactly how patients will react to changes in copayments, both in terms of the new adherence rate and the resulting optimal treatment policy. Future work in this area may include modeling the price elasticity of adherence as an unknown variable within an uncertainty set. Such a formulation lends itself to solutions via robust optimization paired with BLP techniques.

4.7 Conclusion

In conclusion, we have formulated a bilevel optimization model for determining copayment levels faced by a heterogeneous patient population serviced by a single insurance provider. Using the theory of price elasticity, we have modeled the relationship between copayments and the patient's adherence to medication. Each patient class is unique in their characteristics and receives a different clinical benefit from improved adherence to medication. With the knowledge obtained from studying the optimal and personalized treatment policies presented in Chapter III, we developed an iterative nonlinear programming algorithm to derive copayment levels for different patient classes within the population. The class-specific marginal benefit of adherence informs the model's determination of the copayment that should be paid when the insurance provider faces a resource constraint as well as a constraint on the total inequity caused by the copayments. We believe this work is the first to address the problem of improving adherence to personalized and optimal treatment policies. While the formulation was parameterized using patients at risk for cardiovascular disease, the modeling framework is applicable to many chronic diseases where adherence to dynamic treatment plans is a major clinical concern.

CHAPTER V

Conclusion and Future Research

Proper management of patients with chronic diseases can save lives, improve quality of life, and reduce health care costs. This research developed novel methodology and decision support systems to assist physicians, policymakers, and insurance providers with their operational and strategic medical decisions. This work focuses on the use of operations research tools to incorporate elements of personalized medicine in order to better tailor key clinical decisions to the needs of individual patients. The research began with an understanding of how chronic diseases evolve over time and the importance of correctly identifying progressing patients. The research then addressed the development of optimal treatment plans for individual patients in order to improve health outcomes. Finally, using the information learned at the individual patient level, a resource allocation model was developed to exploit the patient level differences and target adherence improving incentives across a heterogeneous patient population.

This body of work has been shown to improve patient outcomes and has contributed to the operations research domain. Chapter II utilized Kalman filtered true state estimates of covariates as data for a generalized estimating equations formulation with a logit link function in order to develop a statistical classification model

with explicit handling of measurement and process noise. This classification model was then embedded within a Kalman filter forecasting model to derive dynamic and personalized monitoring schedules for patients with open-angle glaucoma. Chapter III incorporated conditional value-at-risk to directly account for patient heterogeneity, specifically adherence to medication, in a dynamic programming formulation of optimal hypertension treatment planning. This approach, in combination with a state-space representation of resources, allowed for immediate computation of the marginal benefit of improved adherence and the marginal benefit of increased treatment intensity. Chapter IV developed an iterative nonlinear programming technique to solve a bilevel optimization formulation where the lower level problem was characterized by the dynamic program from Chapter III. The bilevel optimization formulation also addressed the issue of inequity in resource allocations through a constrained social welfare function of the allocations across the heterogeneous patient population.

In conclusion, this research has proposed novel techniques for incorporating the new paradigm of personalized medicine in operations research models of monitoring, treatment, and resource allocation decisions. This research has addressed key clinical questions of interest to physicians, policymakers and insurance providers. Furthermore, this work has generated interest in new areas of research including methodology for capturing medical nuance in treatment planning, as well as theoretical and algorithmic research for solving multilevel optimization models.

5.1 Future Research: Disease Progression Identification

Our prior research in improved statistical classification models combined the systems modeling and noise reduction from Kalman filtering with the marginal analysis of generalized estimating equations to develop a logistic regression function for as-

sessing the probability of experiencing disease progression, given noisy and correlated variables. Logistic regression is a parametric statistical approach which assumes a linear relationship between data and coefficients, i.e. $P(Y = 1 \mid X) = \frac{1}{1+e^{-\beta^T X}}$ where β is a $p \times 1$ vector of coefficients and X is a $p \times 1$ vector of data. Parametric approaches can be limiting in their ability to find the best relationship between the predictor variables X and the outcome variable Y. However, nonparametric approaches (such as random forests) do not suffer from the same restrictions. Future work may combine the Kalman filtering approach with nonparametric approaches to develop semiparametric models for disease progression identification. This approach would retain the noise reduction benefits of Kalman filtering, but would enhance the model through nonparametric estimation. We could then compare the logistic regression approach against the nonparametric models developed to determine which model has the lowest misclassification rate.

5.2 Future Research: Treatment Planning

We consider three extensions to our existing research: (1) growth-mixture modeling for transition probabilities in dynamic programming, (2) optimal treatment initiation strategies for multidrug regimens, and (3) Poisson regression for optimal policy approximation.

The model developed in Chapter III uses a population-based regression model of SBP, HDL and TC in order to forecast cardiovascular risk over time. While the regression model accounts for personalized features such as diabetes and smoking status, theoretically the model assumes the data is drawn from a single population defined by a single set of parameters. To address this issue, multiple-group modeling may be used to develop separate models for each group. However, this approach re-

quires *a priori* knowledge of the groups as well as each patient's membership to one of those groups. To overcome this restriction, we propose to utilize growth-mixture modeling (GMM) to identify multiple unobserved subpopulations [123]. This statistical technique allows for the calculation of longitudinal change both within and across the unobserved subpopulations. GMM maybe used to compute the probability that a patient is a member of each subpopulation based on their observed trajectory. Given repeated measures data and an assumed number of subpopulations, an expectation-maximization or Markov chain Monte Carlo procedure may be employed to iteratively estimate model parameters and compute posterior membership probabilities which provide the maximum likelihood. Then using information criteria like the Akaike Information Criteria (AIC) or Bayesian Information Criteria (BIC), the appropriate final model can be determined.

Given a parameterized GMM for SBP, HDL and TC, the dynamic programming formulation for determining optimal hypertension treatment plans could be greatly improved in the direction of more nuanced personalized medicine. To handle the different subpopulations and each patient's respective membership probabilities, the expected cost-to-go value function V_{t+1} can be computed using the membership probabilities as weights for the transition probabilities for SBP, HDL and TC. Thus, GMM can extend the previously deterministic evolution of these cardiovascular risk factors into stochastic risk factors. Given a sequence of observed SBP, HDL and TC, we will need to dynamically update these membership probabilities (where the membership probabilities at the first decision come directly from the GMM's finalized model parameters). Due to the long planning horizon for chronic disease treatment, the curse of dimensionality quickly becomes problematic if the patient's history of SBP, HDL and TC observations are stored in the state space. As a first step to overcome the

state space explosion, one may assume that membership probabilities are stationary and can be reliably updated via exponential smoothing. Exponential smoothing requires knowledge of only the current observation and the previous smoothed probabilities in order to estimate the current smoothed membership probabilities. Therefore, utilizing GMM and dynamic updating of membership probabilities necessitates including one additional state variable in the dynamic programming formulation.

GMM and dynamic updating of membership probabilities are particularly relevant to medical decision making when there are time-varying treatment benefits and disutilities. In Chapter III, we presented research on the impact of different assumptions on how treatment benefit and disutility depend on the exposure duration to the treatment, e.g. both benefit and disutility increase as the patient takes the medication for longer periods of time. One treatment strategy for incorporating time-varying effects was to base treatment initiation upon a patient's forecasted cardiovascular risk 10 years from today, rather than current protocol based on today's cardiovascular risk. The key feature of the future risk strategy is the ability to identify patient's who will be at high risk for cardiovascular disease events in the distant future so that treatment can be initiated today and the treatment will have had enough time to approach maximum benefit. However, if the forecasting of risk is inaccurate, the truly low-risk patient will accrue all of the compounding treatment disutility but not receive the benefit in the distant future. Therefore, by incorporating GMM and dynamic updating of membership probabilities, the model's forecasts of cardiovascular risk can be greatly enhanced and the decision of when to initiate treatment will be improved.

Moreover, comparisons of treatment initiation strategies under time-varying benefit and disutility were restricted to heuristic risk threshold policies based on either

current or future cardiovascular risk. To optimize the timing of treatment initiation, we propose to update the dynamic programming formulation described above, which includes the GMM and dynamic membership probabilities. In this formulation, the action space may be to wait or initiate treatment, where initiation of treatment implies no future decisions need to be made. In the literature, these stopping time problems formulated as dynamic programs require the computation of the expected cost-to-go value function V_{t+1} when the action is initiate treatment. Typically, this value function is computed (assuming QALYs as the objective function) using a quality of life weight and expected life years remaining when on treatment as estimated from clinical literature. However, under different assumptions of time-varying treatment benefit and disutility, this value function may not be as directly computable. Rather, the dynamic programming formulation may require continued calculations of immediate rewards and cost-to-go value functions using the on-treatment transition probabilities. These transition probabilities are especially challenging to compute when the planning horizon is long, which amplifies the benefits of incorporating GMM and dynamic membership probabilities.

Furthermore, our initial assessment of time-varying benefit and disutility restricted the action space to a single medication initiation. A more robust model would allow for the initiation of several medications over time. To allow for multiple initiations, each with their own functional forms for time-varying benefit and disutility, we could expand the state space to include a medication exposure duration variable η_τ for each possible medication. Assuming a multiplicative treatment benefit and additive treatment disutility when on multiple medications, our approach would then be generalizable to determining optimal initiation strategies for multidrug regimens.

Finally, in the domain of treatment planning, there is great relevance to studying methods for implementing the optimal treatment guidelines in clinical practice. Obtaining the personalized and optimal treatment policy requires solving the dynamic program for each patient a clinician sees. In addition to the barrier of interfacing with the dynamic program, the model may be perceived as a black box to many clinicians which generates distrust of the clinical suggestions. To address both issues, we propose to use policy regression. Policy regression would have the form of a Poisson regression model where the response variable is the number of medications to prescribe and the predictors are patient characteristics including cardiovascular risk, adherence level, and treatment disutility. We propose a three stage approach: (1) training, (2) validation, and (3) testing. The first stage would utilize a large sample of randomized patients and their corresponding optimal treatment policies as inputs to the regression model. Model selection would then be performed using information criteria, such as AIC or BIC, to arrive at the final training model. The output of the final training model will not be guaranteed to be an integer, which may necessitate a rounding rule where the optimal threshold for rounding down is unknown. The second stage might determine the optimal threshold for rounding down by iteratively validating the fit of the Poisson regression predictions against the observed treatment policies. The optimal threshold may then chosen so as to maximize QALYs, which are calculated from a simulation of the rounded predictions across the training data patient population. Finally, the third stage would test the fit and resulting QALYs of the optimal threshold on new randomized patients. This approach provides an easy-to-use model for clinicians that is also easily interpreted. By examining the sign and magnitude of the covariate coefficients in the Poisson regression, clinicians may be better informed as to how the optimal number of medications depends on

patient characteristics, and the coefficients could be used to understand the intuition behind the dynamic program. The Poisson regression model could be easily incorporated into existing electronic health record systems or implemented as a smart phone application.

5.3 Future Research: Copayment Restructuring

In the domain of resource allocation models, we consider the following possible avenues for future research: (1) successive piecewise linear approximation for solving bilevel problems, (2) approximate linear programming via optimal policy and value function regression, and (3) coordination of leaders in hierarchical optimization.

The solution algorithm developed in Chapter IV iteratively solves the lower level MDP for each patient class to obtain the optimal policy given the patient class's adherence. This optimal policy is then given to the upper level nonlinear program to determine locally optimal copayments for each patient class. The limitations of this solution technique is the guarantee of only locally optimal copayments and the lack of scalability to many patient classes. To address both local optimality and scalability, we propose to solve the bilevel problem (BLP) via successive piecewise linear approximations (PLAs). First, we note that the objective function for the BLP can be written as the sum of separable optimal value functions:

$$Z^* = \max_{\mathbf{y}} \left\{ \sum_{t=1}^{T} \sum_{p \in P} \sum_{h \in H} \lambda^t \cdot n_{pht}(\mathbf{y}) \cdot q_{pht} \right\} \tag{5.1}$$

$$= \sum_{p \in P} \sum_{h \in H} n_{ph0} \cdot V_{ph}(\alpha) \tag{5.2}$$

where $n_{pht}(\mathbf{y})$ is the expected number of patients in class p and health state h at time t under decision $\mathbf{y}$ and q_{pht} is the quality of life weight. In the separable objective function (5.2), n_{ph0} is the number of patients in class p and health state h

at the beginning of the planning horizon, and $V_{phT}(s)$ is the optimal value function for patient class p when in health state h with adherence rate α at the beginning of the planning horizon. This optimal value function comes from solving the MDP (see equation (3.19)) for patient class p when starting in health state h with adherence α that depends on the patient class's copayment level. Given this separability, if we can approximate each optimal function linearly, we can reduce the nonlinearity of the BLP.

To approximate the optimal value function, we can use piecewise linear functions $\hat{f}_{ph}(\alpha)$ to approximate $V_{ph}(\alpha)$ as piecewise linear in α (the outcome variable associated with our copayment coverage decision variable y_{ph}). PLAs connect points (knots) of the original function with lines of the form $y = mx + b$ (standard linear form denoting an outcome y as the product of a slope m and parameter x plus an intercept b) to approximate the entire function with a series of linear functions [79]. PLAs have the following form:

$$\hat{f}(\alpha) = f(x_k) + (\alpha - x_k)\frac{f(x_{k+1}) - f(x_k)}{x_{k+1} - x_k} \tag{5.3}$$

where $x_k, k = 1, \ldots, K$ are knots and f is the original function. Typically, one selects the set of knots x_k beforehand, approximates the function f, and uses the PLA $\hat{f}$ in the optimization problem [131]. We can then write our objective function as:

$$Z^* = \sum_{p \in P} \sum_{h \in H} n_{ph0} \cdot V_{ph}(\alpha) \tag{5.4}$$

$$= \sum_{p \in P} \sum_{h \in H} n_{ph0} \cdot \hat{f}_{ph}(\alpha) \tag{5.5}$$

Rather than solve the BLP with a single set of knots, we propose a successive approximation scheme that updates the set of knots to improve the fit of the approximations $\hat{f}_{ph}(\alpha)$. Our proposed algorithm is described below:

1. For each $p \in P$, solve the MDP for the current adherence vector α^i_{pht}, the maximum adherence $\alpha^M_{pht} = \alpha^i_{pht} + \delta^\alpha_{pht} \cdot (1 - y^i_{ph})$ and the minimum adherence $\alpha^m_{pht} = \alpha^i_{pht} + \delta^\alpha_{pht} \cdot (0 - y^i_{ph})$. Obtain the optimal policies π^* and the optimal value functions $V_{ph}(\alpha)$.

2. For each combination of p and h, fit piecewise linear approximations (piecewise linear in α) to $V_{ph}(\alpha)$ at the knots: $\mathbf{x}_{ph} = (\alpha^m_{pht}, \alpha^i_{pht}, \alpha^M_{pht})$.

3. Using the PLAs, solve the BLP with objective defined by equation (5.5) to find $\mathbf{y}^*$.

4. For each $p \in P$, solve the MDP at the solution y^*_{ph}. Update the objective function value Z^*. Add $\alpha_{pht} = \max\{0, \min\{1, \alpha^i_{pht} + \delta^\alpha_{pht} \cdot (y_{ph} - y^i_{ph})\}\}$ to the set of knots $\mathbf{x}_{ph}$. Solve the MDP at the new knot and update the piecewise linear approximations.

5. Repeat Steps 3 and 4 until Z^* is within convergence requirement.

The successive PLA algorithm initializes the set of knots for PLA as the minimum, current and maximum adherence for patient class p, given their price elasticity of adherence δ^α_{pht}. These initial knots cover the full range of possible values for the optimal value function $V_{ph}(\alpha)$. Given the PLA under the initial set of knots, we may then solve the BLP to obtain a globally optimal solution to the approximate objective function. Next, we could add the adherence corresponding to the optimal solution to the set of knots. The MDP would then be solved at the new knot, the PLA would be updated, and the BLP would be solved again. We may continue this process until the objective function value Z^* has not sufficiently changed after expanding the set of knots in the PLAs.

This algorithm has the following features: (1) monotonicity of $V_{ph}(\alpha)$ in α leads to improving optimal BLP solutions and objective function value Z^*, (2) the objective function value at the BLP optimal solution is exactly equal to the original nonlinear objective function value, (3) under concave BLP constraints, the feasibility of the BLP solution for the PLAs implies an optimal solution to the original BLP formulation, and (4) the piecewise linearity of the objective function may improve the scalability of the BLP to include more patient classes.

We propose to use this successive PLA algorithm to solve the BLP formulation in Chapter IV and compare the performance of the two solution techniques. Given the features described above, we believe the successive PLA algorithm should yield faster and better solutions on problems of larger size.

Another approximation approach for solving the original BLP formulation is to use approximate linear programming (ALP) by approximating the objective function and constraints using basis functions [41]. A linear basis function approximation for the optimal value function $V(s)$ has the following form:

$$V(s) \approx \hat{v}(s) = \mathbf{p}^T \beta(s) \tag{5.6}$$

where β is a vector of basis functions for state s and $\mathbf{p}$ is a vector of basis weights. The goal of $\hat{v}(s)$ is to find the weights $\mathbf{p}$ that minimizes the error between $V(s)$ and $\hat{v}(s)$. From a statistics perspective, finding the weights $\mathbf{p}$ is equivalent to determining regression coefficients β in the regression function $Y = \beta^T X$. Therefore, we can approximate the optimal value function $V(s)$ via linear regression. Similarly, we could approximate the optimal policy π^* from the MDP described in Chapter III using linear regression. These linear regression models can be parameterized using data from randomly sampled patients and their respective MDP solutions and policies. Pairing linear representations of the objective function and optimal treatment plans

with linearizations of the budget and equity constraints, the original BLP can be reformulated as a linear program (LP). The LP has the property of global optimality and very efficient solution methods [40]. Given linear approximations, we can rewrite the original BLP as:

$$\max_{\mathbf{y},\hat{\boldsymbol{\pi}}(\mathbf{y})} \quad \mathbf{p}_1^T \beta_1(\mathbf{y}, \hat{\boldsymbol{\pi}}(\mathbf{y})) \tag{5.7}$$

$$\text{s.t.} \quad \hat{G}_i(\mathbf{y}, \hat{\boldsymbol{\pi}}(\mathbf{y})) \leq 0 \quad \forall i \tag{5.8}$$

$$\hat{\pi}_p(\mathbf{y}) = \mathbf{p}_2^T \beta_2(\mathbf{y}) \quad \forall p \tag{5.9}$$

$$\hat{g}_{pj}(\mathbf{y}, \hat{\pi}_p) \leq 0 \quad \forall p, j \tag{5.10}$$

where $\mathbf{p}_1$ and $\mathbf{p}_2$ are weights for the leader's objective function and the follower's optimal treatment policy, respectively. Similarly, β_1 and β_2 are basis function vectors for the leader's objective function and the follower's optimal treatment policy, respectively. And $\hat{G}_i$ and $\hat{g}_{pj}$ are linear approximations to the leader's and follower's constraints, respectively.

The benefits of the ALP approach over PLA is the solution speed. Once parameterized, the ALP approach does not require repeatedly solving the MDP for each patient class $p \in P$. Furthermore, the efficiency of linear programming solution techniques, like simplex, imply that the ALP approach may scale better with the problem size than the successive PLA approach. However, because of the degree of approximations necessary to formulate the BLP as a LP, the optimal solution to the ALP may not be sufficiently close to the optimal solution to the BLP. Futher analysis and implementation studies are necessary to properly evaluate the trade-offs between the two approximation techniques.

Lastly, we consider modeling multiple insurance providers (i.e. leaders in the BLP), where each services their own heterogeneous patient population. In our prior

work, we assumed a single insurance provider. Under a free market with multiple private insurance providers, the patients are able to change insurance providers during the open enrollment season [152]. Assuming rational patients who seek to minimize their costs and insurance providers who are otherwise equal in attractiveness, if an insurance provider decreases copayments for a patient class (e.g. diabetics), patients of that class will migrate to that insurance provider during open enrollment. This may leave the insurance provider who decreased copayments with an influx of high risk patients (they are prioritized in the QALY-maximizing BLP formulation) which increases their costs. This leaves the other insurance providers with a generally healthier patient population and lower costs. Therefore, there is a first-mover disadvantage to reducing copayments when there are multiple private insurance providers servicing a patient population who can switch providers. Given these problems, it may be beneficial for the government to allocate resources to the insurance providers to incentivize the reduction of copayments. Adding government intervention to the system adds a new optimization level to the hierarchical model, creating a trilevel optimization problem where the highest level is the government, the next level is the insurance providers, and the lower level is the patients.

The existence of multiple insurance providers occurs in the private insurance market where the insurance provider may be interested in minimizing expected costs rather than maximizing expected total health of its patient population. However, we may assume the government wants to maximize the expected health of the population when allocating resources to encourage copayment reform. Therefore, the

trilevel optimization problem may have the following form:

$$\max_{\mathbf{x}} \quad H(\mathbf{x}, \mathbf{y}(\mathbf{x}), \boldsymbol{\pi}(\mathbf{y}(\mathbf{x}))) \tag{5.11}$$

$$\text{s.t.} \quad h_l(\mathbf{x}, \mathbf{y}(\mathbf{x}), \boldsymbol{\pi}(\mathbf{y}(\mathbf{x}))) \leq 0 \quad \forall l \tag{5.12}$$

$$\mathbf{y}_r(\mathbf{x}) = \arg\min_{\mathbf{y}} F_r(\mathbf{y}(\mathbf{x}), \boldsymbol{\pi}(\mathbf{y}(\mathbf{x}))) \quad \forall r \tag{5.13}$$

$$G_{ri}(\mathbf{y}_r(\mathbf{x}), \boldsymbol{\pi}(\mathbf{y}_r(\mathbf{x}))) \leq 0 \quad \forall r, i \tag{5.14}$$

$$\pi_{pr}(\mathbf{y}_r(\mathbf{x})) = \arg\max_{\pi} f_p(\mathbf{y}_r(\mathbf{x}), \pi) \quad \forall p, r \tag{5.15}$$

$$g_{prj}(\mathbf{y}_r(\mathbf{x}), \pi_p) \leq 0 \quad \forall p, r, j \tag{5.16}$$

where equation (5.11) captures the government's interest in maximizing total health of the population by allocating resource vector $\mathbf{x}$ across each insurance provider $r \in R$. Equation (5.12) denotes each of the $l \in L$ constraints faced by the goverment, including budgetary constraints on the allocations to the insurance providers. Equations (5.13) now reflects the insurance providers' desire to minimize their expected total costs F_r, which depend on the government's allocation $\mathbf{x}$ as well as the reactions of its patients to the copayment level $\mathbf{y}_r$. Each private insurance provider will have its own budget and equity constraints, denoted by equation (5.14). The optimal treatment policy now depends on the copayment coverage set by their insurance provider as described in equation (5.15). In addition the minimum SBP thresholds, each patient may switch insurance providers in accordance with equation (5.16).

With this trilevel optimization (TLP) formulation, we can study the structure of a coordination strategy set by the government that would lead private insurance providers to reduce copayments so that patient health can be improved via increased adherence to medication. Furthermore, we can analyze how differences in budgets and equity constraints amongst the private insurance providers influence the TLP resource allocation and the corresponding optimal value. Numerically, we could con-

sider how state governments can use incentives to coordinate copayment reductions amongst a set of insurance providers servicing a particular geographical region.

The extra level of decisions makes this hierarchical model more challenging to solve than the BLP described in Chapter IV. In fact, the BLP is embedded within the TLP. One solution technique may be to solve the BLP under different government allocations and use this series of solutions to approximate equations (5.13) through (5.16). By approximating the outcome of government allocations to each private insurance provider, we may be able to reduce the hierarhical complexity to a single level mathematical program. Depending on the convexity of H and the shape of the approximation functions for equations (5.13) through (5.16), we may be able to solve the approximate problem using existing techniques.

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] ADA. American diabetes assocation: Statistics about diabetes. *http://www.diabetes.org/diabetes-basics/statistics/*, June 2014.

[2] Jae-Hyeon Ahn and John C Hornberger. Involving patients in the cadaveric kidney transplant allocation process: A decision-theoretic perspective. *Management Science*, 42(5):629–641, 1996.

[3] Oguzhan Alagoz, Heather Hsu, Andrew J Schaefer, and Mark S Roberts. Markov decision processes: a tool for sequential decision making under uncertainty. *Medical Decision Making*, 30(4):474–483, 2010.

[4] Oguzhan Alagoz, Lisa M Maillart, Andrew J Schaefer, and Mark S Roberts. The optimal timing of living-donor liver transplantation. *Management Science*, 50(10):1420–1430, 2004.

[5] Oguzhan Alagoz, Lisa M Maillart, Andrew J Schaefer, and Mark S Roberts. Determining the acceptance of cadaveric livers using an implicit model of the waiting list. *Operations Research*, 55(1):24–36, 2007.

[6] P.S. Albert and L.M. McShane. A generalized estimating equations approach for spatially correlated binary data: applications to the analysis of neuroimaging data. *Biometrics*, 51(2):627–638, 1995.

[7] Gerard F Anderson. *Chronic care: making the case for ongoing care*. Robert Wood Johnson Foundation, 2010.

[8] Keaven M Anderson, Patricia M Odell, Peter WF Wilson, and William B Kannel. Cardiovascular disease risk profiles. *American heart journal*, 121(1):293–298, 1991.

[9] Elizabeth Arias. United states life tables, 2004. *National vital statistics reports*, 56(9):1–40, 2007.

[10] Elizabeth Arias. United states life tables, 2007. *National vital statistics reports: from the Centers for Disease Control and Prevention, National Center for Health Statistics, National Vital Statistics System*, 59(9):1, 2011.

[11] American Lung Association. Smoking and older adults.

[12] Turgay Ayer, Oguzhan Alagoz, and Natasha K Stout. Pomdp approach to personalize mammography screening decisions. *Operations Research*, 60(5):1019–1034, 2012.

[13] Mehmet US Ayvaci, Oguzhan Alagoz, and Elizabeth S Burnside. The effect of budgetary restrictions on breast cancer diagnostic decisions. *Manufacturing & Service Operations Management*, 14(4):600–617, 2012.

[14] Sripal Bangalore, Jie Qin, Sarah Sloan, Sabina A Murphy, Christopher P Cannon, et al. What is the optimal blood pressure in patients after acute coronary syndromes? relationship of blood pressure and cardiovascular events in the pravastatin or atorvastatin evaluation and infection therapy–thrombolysis in myocardial infarction (prove it-timi) 22 trial. *Circulation*, 122(21):2142–2151, 2010.

[15] David R Bangsberg, Travis C Porco, C Kagay, Edwin D Charlebois, Steven G Deeks, David Guzman, Richard Clark, and Andrew Moss. Modeling the hiv protease inhibitor adherence-resistance curve by use of empirically derived estimates. *Journal of Infectious Diseases*, 190(1):162–165, 2004.

[16] T Başar and R Srikant. A stackelberg network game with a large number of followers. *Journal of optimization theory and applications*, 115(3):479–490, 2002.

[17] B. Bengtsson, V.M. Patella, and A. Heijl. Prediction of glaucomatous visual field loss by extrapolation of linear trends. *Archives of ophthalmology*, 127(12):1610, 2009.

[18] Erik JAJ Beune, Eric P Moll van Charante, Leo Beem, Jacob Mohrs, Charles O Agyemang, Gbenga Ogedegbe, and Joke A Haafkens. Culturally adapted hypertension education (cahe) to improve blood pressure control and treatment adherence in patients of african origin with uncontrolled hypertension: Cluster-randomized trial. *PloS one*, 9(3):e90103, 2014.

[19] Kang Boda and Jerzy A Filar. Time consistent dynamic risk measures. *Mathematical Methods of Operations Research*, 63(1):169–186, 2006.

[20] Vivek Borkar and Rahul Jain. Risk-constrained markov decision processes. In *Decision and Control (CDC), 2010 49th IEEE Conference on*, pages 2664–2669. IEEE, 2010.

[21] Henrik Brønnum-Hansen, Torben Jørgensen, Michael Davidsen, Mette Madsen, Merete Osler, Lars Ulrik Gerdes, and Marianne Schroll. Survival and cause of death after myocardial infarction:: The danish monica study. *Journal of clinical epidemiology*, 54(12):1244–1250, 2001.

[22] O. Brovko, DM Wiberg, L. Arena, and JW Bellville. The extended kalman filter as a pulmonary blood flow estim [combining dot above] ator. *Automatica*, 17(1):213–220, 1981.

[23] Dieu Tien Bui, Tran Anh Tuan, Harald Klempe, Biswajeet Pradhan, and Inge Revhaug. Spatial prediction models for shallow landslide hazards: a comparative assessment of the efficacy of support vector machines, artificial neural networks, kernel logistic regression, and logistic model tree. *Landslides*, pages 1–18, 2015.

[24] Barry A Bunting, Benjamin H Smith, and Susan E Sutherland. The asheville project: clinical and economic outcomes of a community-based long-term medication therapy management program for hypertension and dyslipidemia. *Journal of the American Pharmacists Association: JAPhA*, 48(1), 2008.

[25] John Burn, Martin Dennis, John Bamford, Peter Sandercock, Derick Wade, and Charles Warlow. Long-term risk of recurrent stroke after a first-ever stroke. the oxfordshire community stroke project. *Stroke*, 25(2):333–337, 1994.

[26] Mary K Catlin, John A Poisal, and Cathy A Cowan. Out-of-pocket health care expenditures, by insurance status, 2007–10. *Health Affairs*, 34(1):111–116, 2015.

[27] CDC. Center for disease control and prevention: Smoking and tobacco use. *http://www.cdc.gov/tobacco/data_statistics/fact_sheets/fast_facts/*, April 2014.

[28] Jenny C Chang, Eric C Wooten, Anna Tsimelzon, Susan G Hilsenbeck, M Gutierrez, Richard Elledge, Syed Mohsin, C Kent Osborne, Gary C Chamness, D Craig Allred, et al. Gene expression profiling for the prediction of therapeutic response to docetaxel in patients with breast cancer. *The Lancet*, 362(9381):362–369, 2003.

[29] Michael E Chernew, William E Encinosa, and Richard A Hirth. Optimal health insurance: the case of observable, severe illness. *Journal of Health Economics*, 19(5):585–609, 2000.

[30] Michael E Chernew, Mayur R Shah, Arnold Wegh, Stephen N Rosenberg, Iver A Juster, Allison B Rosen, Michael C Sokol, Kristina Yu-Isenberg, and A Mark Fendrick. Impact of decreasing copayments on medication adherence within a disease management environment. *Health Affairs*, 27(1):103–112, 2008.

[31] Jagpreet Chhatwal, Oguzhan Alagoz, and Elizabeth S Burnside. Optimal breast biopsy decision-making based on mammographic features and demographic factors. *Operations research*, 58(6):1577–1591, 2010.

[32] Aram V Chobanian, George L Bakris, Henry R Black, William C Cushman, Lee A Green, Joseph L Izzo, Daniel W Jones, Barry J Materson, Suzanne Oparil, Jackson T Wright, et al. Seventh report of the joint national committee on prevention, detection, evaluation, and treatment of high blood pressure. *Hypertension*, 42(6):1206–1252, 2003.

[33] Niteesh K Choudhry, Jerry Avorn, Elliott M Antman, Sebastian Schneeweiss, and William H Shrank. Should patients receive secondary prevention medications for free after a myocardial infarction? an economic analysis. *Health Affairs*, 26(1):186–194, 2007.

[34] Niteesh K Choudhry, Michael A Fischer, Jerry Avorn, Sebastian Schneeweiss, Daniel H Solomon, Christine Berman, Saira Jan, Jun Liu, Joyce Lii, M Alan Brookhart, et al. At pitney bowes, value-based insurance design cut copayments and increased drug adherence. *Health Affairs*, 29(11):1995–2001, 2010.

[35] Niteesh K Choudhry, Amanda R Patrick, Elliott M Antman, Jerry Avorn, and William H Shrank. Cost-effectiveness of providing full drug coverage to increase medication adherence in post–myocardial infarction medicare beneficiaries. *Circulation*, 117(10):1261–1268, 2008.

[36] Karl Claxton, Mike Paulden, Hugh Gravelle, Werner Brouwer, and Anthony J Culyer. Discounting and decision making in the economic evaluation of health-care technologies. *Health economics*, 20(1):2–15, 2011.

[37] Paul Contoyannis, Jeremiah Hurley, Paul Grootendorst, Sung-Hee Jeon, and Robyn Tamblyn. Estimating the price elasticity of expenditure for prescription drugs in the presence of nonlinear price schedules: an illustration from quebec, canada. *Health Economics*, 14(9):909–923, 2005.

[38] Jeffrey A Cutler, Paul D Sorlie, Michael Wolz, Thomas Thom, Larry E Fields, and Edward J Roccella. Trends in hypertension prevalence, awareness, treatment, and control rates in united states adults between 1988–1994 and 1999–2004. *Hypertension*, 52(5):818–827, 2008.

[39] Ralph B D'Agostino, Mason W Russell, Daniel M Huse, R Curtis Ellison, Halit Silbershatz, Peter WF Wilson, and Stuart C Hartz. Primary and subsequent coronary risk appraisal: new results from the framingham study. *American heart journal*, 139(2):272–281, 2000.

[40] George Bernard Dantzig. *Linear programming and extensions*. Princeton university press, 1998.

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

[42] Luca Degli Esposti and Giorgia Valpiani. Pharmacoeconomic burden of undertreating hypertension. *Pharmacoeconomics*, 22(14):907–928, 2004.

[43] Stephan Dempe and Alain B Zemkoho. Kkt reformulation and necessary conditions for optimality in nonsmooth bilevel optimization. *SIAM Journal on Optimization*, 24(4):1639–1669, 2014.

[44] Cheryl Dennison-Himmelfarb, Joel Handler, Daniel T Lackland, Michael L LeFevre, Thomas D MacKenzie, Olugbenga Ogedegbe, Sidney C Smith Jr, Laura P Svetkey, Sandra J Taler, Raymond R Townsend, et al. 2014 evidence-based guideline for the management of high blood pressure in adults report from the panel members appointed to the eighth joint national committee (jnc 8). *Journal of the American Medical Association*, 2013.

[45] Brian T Denton, Murat Kurt, Nilay D Shah, Sandra C Bryant, and Steven A Smith. Optimizing the start time of statin therapy for patients with diabetes. *Medical Decision Making*, 29(3):351–367, 2009.

[46] V.T. Diaz-Aleman, A. Anton, M.G. de la Rosa, ZK Johnson, S. McLeod, and A. Azuara-Blanco. Detection of visual-field deterioration by Glaucoma Progression Analysis and Threshold Noiseless Trend programs. *British Journal of Ophthalmology*, 93(3):322, 2009.

[47] B.N. Doebbeling, M.B. Edmond, C.S. Davis, J.R. Woodin, and R.R. Zeitler. Influenza Vaccination of Health Care Workers: Evaluation of Factors That Are Important in Acceptance* 1. *Preventive medicine*, 26(1):68–77, 1997.

[48] William E Evans and Mary V Relling. Moving towards individualized medicine with pharmacogenomics. *Nature*, 429(6990):464–468, 2004.

[49] Daniel M Faissol, Paul M Griffin, and Julie L Swann. Timing of testing and treatment of hepatitis c and other diseases. In *INFORMS international meeting*, pages 8–11, 2007.

[50] Marianna Fontana, Perviz Asaria, Michela Moraldo, Judith Finegold, Khalil Hassanally, Charlotte H Manisty, and Darrel P Francis. Patient-accessible tool for shared decision making in cardiovascular primary prevention balancing longevity benefits against medication disutility. *Circulation*, 129(24):2539–2546, 2014.

[51] National Center for Chronic Disease Prevention and Health Promotion. Chronic diseases and health promotion.

[52] Centers for Disease Control, Prevention, et al. National diabetes statistics report: estimates of diabetes and its burden in the united states, 2014. *Atlanta, GA: US Department of Health and Human Services*, 2014.

[53] Allen M Fremont, Paul P Lee, Carol M Mangione, Kanika Kapur, John L Adams, Steven L Wickstrom, and José J Escarce. Patterns of care for open-angle glaucoma in managed care. *Archives of ophthalmology*, 121(6):777–783, 2003.

[54] DS Friedman, RC Wolfs, BJ O'colmain, BE Klein, HR Taylor, S. West, MC Leske, P. Mitchell, N. Congdon, and J. Kempen. Prevalence of open-angle glaucoma among adults in the United States. *Archives of ophthalmology*, 122(4):532, 2004.

[55] Dennis G Fryback, Erik J Dasbach, Ronald Klein, Barbara EK Klein, Norma Dorn, Kathy Peterson, and Patrica A Martin. The beaver dam health outcomes study initial catalog of health-state quality factors. *Medical Decision Making*, 13(2):89–102, 1993.

[56] Prashant C Fuloria and Stefanos A Zenios. Outcomes-adjusted reimbursement in a healthcare delivery system. *Management Science*, 47(6):735–751, 2001.

[57] Z. Ghahramani and G.E. Hinton. Parameter estimation for linear dynamical systems. *University of Toronto technical report CRG-TR-96-2*, 6, 1996.

[58] Archis Ghate and Minsum Kim. Dynamic programming for response-guided therapy. 2014.

[59] Ihab Hajjar and Theodore A Kotchen. Trends in prevalence, awareness, treatment, and control of hypertension in the united states, 1988-2000. *JAMA: the journal of the American Medical Association*, 290(2):199–206, 2003.

[60] Ulrich Halekoh, Søren Højsgaard, and Jun Yan. The r package geepack for generalized estimating equations. *Journal of Statistical Software*, 15(2):1–11, 2006.

[61] Margaret A Hamburg and Francis S Collins. The path to personalized medicine. *New England Journal of Medicine*, 363(4):301–304, 2010.

[62] Pierre Hansen, Brigitte Jaumard, and Gilles Savard. New branch-and-bound rules for linear bilevel programming. *SIAM Journal on Scientific and Statistical Computing*, 13(5):1194–1217, 1992.

[63] Milos Hauskrecht and Hamish Fraser. Planning treatment of ischemic heart disease with partially observable markov decision processes. *Artificial Intelligence in Medicine*, 18(3):221–244, 2000.

[64] J. Helm, M.S. Lavieri, M. Van Oyen, J. Stein, and D. Musch. Dynamic forecasting and control algorithms of glaucoma progression for clinician decision support. Paper Under Review.

[65] Hanane Hemi, Jamel Ghouili, and Ahmed Cheriti. Combination of markov chain and optimal control solved by pontryagins minimum principle for a fuel cell/supercapacitor vehicle. *Energy Conversion and Management*, 91:387–393, 2015.

[66] Melonie Heron. Deaths: leading causes for 2004. *National vital statistics reports*, 56(5):1–96, 2007.

[67] David U Himmelstein and Steffie Woolhandler. Cost control in a parallel universe: Medicare spending in the united states and canada. *Archives of internal medicine*, 172(22):1764–1766, 2012.

[68] P Michael Ho, Anne Lambert-Kerzner, Evan P Carey, Ibrahim E Fahdi, Chris L Bryson, S Dee Melnyk, Hayden B Bosworth, Tiffany Radcliff, Ryan Davis, Howard Mun, et al. Multifaceted intervention to improve medication adherence and secondary prevention measures after acute coronary syndrome hospital discharge: a randomized clinical trial. *JAMA internal medicine*, 174(2):186–193, 2014.

[69] E. Hodapp, R.K. Parrish II, D.R. Anderson, and T.W. Perkins. *Clinical decisions in glaucoma*. Mosby, St. Louis, MO, 1993.

[70] Chuanpu Hu, William S Lovejoy, and Steven L Shafer. Comparison of some suboptimal control policies in medical drug therapy. *Operations Research*, 44(5):696–709, 1996.

[71] Robert Hutchins, Anthony J Viera, Stacey L Sheridan, and Michael P Pignone. Quantifying the utility of taking pills for cardiovascular prevention. *Circulation: Cardiovascular Quality and Outcomes*, pages CIRCOUTCOMES–114, 2015.

[72] Houyuan Jiang, Zhan Pang, and Sergei Savin. Performance-based contracts for outpatient medical services. *Manufacturing & Service Operations Management*, 14(4):654–669, 2012.

[73] R.E. Kalman et al. A new approach to linear filtering and prediction problems. *Journal of basic Engineering*, 82(1):35–45, 1960.

[74] Leslie A Kenna, Line Labbé, Jeffrey S Barrett, and Marc Pfister. Modeling and simulation of adherence: approaches and applications in therapeutics. *The AAPS journal*, 7(2):E390–E407, 2005.

[75] Gregory James King. *Essays on Service and Health Care Operations*. PhD thesis, The University of Michigan, 2013.

[76] Eser Kırkızlar, Daniel M Faissol, Paul M Griffin, and Julie L Swann. Timing of testing and treatment for asymptomatic diseases. *Mathematical biosciences*, 226(1):28–37, 2010.

[77] Kenneth D Kochanek, Jiaquan Xu, Sherry L Murphy, Arialdi M Minino, and Hsiang-Ching kung. Deaths: final data for 2009. *National vital statistics reports*, 60(3), 2011.

[78] Huiberdina L Koek, Agnes de Bruin, Fred Gast, Evelien Gevers, Jan WPF Kardaun, Johannes B Reitsma, Diederick E Grobbee, and Michiel L Bots. Short-and long-term prognosis after acute myocardial infarction in men versus women. *The American journal of cardiology*, 98(8):993–999, 2006.

[79] Spyros Kontogiorgis. Practical piecewise-linear approximation for monotropic optimization. *INFORMS Journal on Computing*, 12(4):324–340, 2000.

[80] Amanda E Kowalski. Censored quantile instrumental variable estimates of the price elasticity of expenditure on medical care. Technical report, National Bureau of Economic Research, 2009.

[81] Jennifer E Kreke, Matthew D Bailey, Andrew J Schaefer, Derek C Angus, and Mark S Roberts. Modeling hospital discharge policies for patients with pneumonia-related sepsis. *IIE Transactions*, 40(9):853–860, 2008.

[82] Pavlo Krokhmal, Jonas Palmquist, and Stanislav Uryasev. Portfolio optimization with conditional value-at-risk objective and constraints. *Journal of Risk*, 4:43–68, 2002.

[83] Marie Krousel-Wood, Amanda Hyre, Paul Muntner, and Donald Morisky. Methods to improve medication adherence in patients with hypertension: current status and future directions. *Current opinion in cardiology*, 20(4):296–300, 2005.

[84] Murat Kurt, Brian T Denton, Andrew J Schaefer, Nilay D Shah, and Steven A Smith. The structure of optimal statin initiation policies for patients with type 2 diabetes. *IIE Transactions on Healthcare Systems Engineering*, 1(1):49–65, 2011.

[85] M. Kuure-Kinsey, C.C. Palerm, and B.W. Bequette. A dual-rate kalman filter for continuous glucose monitoring. In *Engineering in Medicine and Biology Society, 2006. EMBS'06. 28th Annual International Conference of the IEEE*, pages 63–66. IEEE, 2006.

[86] Darius Lakdawalla and Neeraj Sood. Innovation and the welfare effects of public drug insurance. *Journal of public economics*, 93(3):541–548, 2009.

[87] M.S. Lavieri, M.L. Puterman, S. Tyldesley, and W.J. Morris. When to treat prostate cancer patients based on their psa dynamics. 2012.

[88] MR Law, JK Morris, and NJ Wald. Use of blood pressure lowering drugs in the prevention of cardiovascular disease: meta-analysis of 147 randomised trials in the context of expectations from prospective epidemiological studies. *BMJ: British Medical Journal*, 338, 2009.

[89] Donald KK Lee and Stefanos A Zenios. An evidence-based incentive system for medicare's end-stage renal disease program. *Management Science*, 58(6):1092–1105, 2012.

[90] P.P. Lee, J.W. Walt, L.C. Rosenblatt, L.R. Siegartel, and L.S. Stern. Association between intraocular pressure variation and glaucoma progression: data from a United States chart review. *American journal of ophthalmology*, 144(6):901–907, 2007.

[91] L.A. Lenert, A. Sturley, and M. Rupnow. Toward improved methods for measurement of utility: automated repair of errors in elicitations. *Medical decision making*, 23(1):67, 2003.

[92] Hecheng Li and Yuping Wang. A hybrid genetic algorithm for solving nonlinear bilevel programming problems based on the simplex method. In *Natural Computation, 2007. ICNC 2007. Third International Conference on*, volume 4, pages 91–95. IEEE, 2007.

[93] Y Li, BI Graubard, P Huang, and JL Gastwirth. Extension of the peters–belson method to estimate health disparities among multiple groups using logistic regression with survey data. *Statistics in medicine*, 34(4):595–612, 2015.

[94] Ting Liu and Ching-to Albert Ma. Health insurance, treatment plan, and delegation to altruistic physician. *Journal of Economic Behavior & Organization*, 85:79–96, 2013.

[95] Donald Lloyd-Jones, Robert Adams, Mercedes Carnethon, Giovanni De Simone, T Bruce Ferguson, Katherine Flegal, Earl Ford, Karen Furie, Alan Go, Kurt Greenlund, et al. Heart disease and stroke statistics2009 update a report from the american heart association statistics committee and stroke statistics subcommittee. *Circulation*, 119(3):e21–e181, 2009.

[96] Reza Mahjoub, Fredrik Odegaard, and Gregory S Zaric. Health-based pharmaceutical pay-for-performance risk-sharing agreements. *Journal of the Operational Research Society*, 2013.

[97] Bikaramjit S Mann, Lianne Barnieh, Karen Tang, David JT Campbell, Fiona Clement, Brenda Hemmelgarn, Marcello Tonelli, Diane Lorenzetti, and Braden J Manns. Association between drug insurance cost sharing strategies and outcomes in patients with chronic diseases: A systematic review. *PloS one*, 9(3):e89168, 2014.

[98] Steve L Mansberger, Felipe A Medeiros, and Mae Gordon. Diagnostic tools for calculation of glaucoma risk. *Survey of ophthalmology*, 53(6):S11–S16, 2008.

[99] Steven L Mansberger and George A Cioffi. The probability of glaucoma from ocular hypertension determined by ophthalmologists in comparison to a risk calculator. *Journal of glaucoma*, 15(5):426–431, 2006.

[100] Steven I Marcus, Emmanual Fernández-Gaucherand, Daniel Hernández-Hernandez, Stefano Coraluppi, and Pedram Fard. Risk sensitive markov decision processes. *Progress in Systems and Control Theory*, 22:263–280, 1997.

[101] T. Marshall. Misleading measurements: modeling the effects of blood pressure misclassification in a united states population. *Medical decision making*, 26(6):624, 2006.

[102] Jennifer E Mason, Brian T Denton, Nilay D Shah, and Steven A Smith. Optimizing the simultaneous management of blood pressure and cholesterol for type 2 diabetes patients. *European Journal of Operations Research*, 2013.

[103] Jennifer E Mason, Darin A England, Brian T Denton, Steven A Smith, Murat Kurt, and Nilay D Shah. Optimizing statin treatment decisions for diabetes patients in the presence of uncertain future adherence. *Medical Decision Making*, 32(1):154–166, 2012.

[104] Giampiero Mazzaglia, Ettore Ambrosioni, Marianna Alacqua, Alessandro Filippi, Emiliano Sessa, Vincenzo Immordino, Claudio Borghi, Ovidio Brignoli, Achille P Caputi, Claudio Cricelli, et al. Adherence to antihypertensive medications and cardiovascular morbidity among newly diagnosed hypertensive patients. *Circulation*, 120(16):1598–1605, 2009.

[105] J. McNames and M. Aboy. Statistical modeling of cardiovascular signals and parameter estimation based on the extended kalman filter. *Biomedical Engineering, IEEE Transactions on*, 55(1):119–129, 2008.

[106] AI McNaught, R.A. Hitchings, DP Crabb, and FW Fitzke. Modelling series of visual fields to detect progression in normal-tension glaucoma. *Graefe's Archive for Clinical and Experimental Ophthalmology*, 233(12):750–755, 1995.

[107] Sherry L Murphy, Jiaquan Xu, and Kenneth D Kochanek. Deaths: final data for 2010. *National vital statistics reports: from the Centers for Disease Control and Prevention, National Center for Health Statistics, National Vital Statistics System*, 61(4):1–117, 2013.

[108] David C Musch, Brenda W Gillespie, Paul R Lichter, Leslie M Niziol, Nancy K Janz, and CIGTS Study Investigators. Visual field progression in the collaborative initial glaucoma treatment study: the impact of treatment and other baseline factors. *Ophthalmology*, 116(2):200–207, 2009.

[109] D.C. Musch, B.W. Gillespie, L.M. Niziol, L.F. Cashwell, and P.R. Lichter. Factors associated with intraocular pressure before and during 9 years of treatment in the Collaborative Initial Glaucoma Treatment Study. *Ophthalmology*, 115(6):927–933, 2008.

[110] Stéphane Mussard, Françoise Seyte, and Michel Terraza. Decomposition of gini and the generalized entropy inequality measures. *Economics Bulletin*, 4(7):1–6, 2003.

[111] NIAAA. National institute on alcohol abuse and alcoholism: Alcohol facts and statistics. *http://www.niaaa.nih.gov/alcohol-health/overview-alcohol-consumption/alcohol-facts-and-statistics*, July 2014.

[112] David Nolan, Silvana Gaudieri, Mina John, and Simon Mallal. Impact of host genetics on hiv disease progression and treatment: new conflicts on an ancient battleground. *Aids*, 18(9):1231–1240, 2004.

[113] K. Nouri-Mahdavi, J. Caprioli, A.L. Coleman, D. Hoffman, and D. Gaasterland. Pointwise linear regression for evaluation of visual field outcomes and comparison with the advanced glaucoma intervention study methods. *Archives of ophthalmology*, 123(2):193, 2005.

[114] Jonathan Oberlander. Unfinished journeya century of health care reform in the united states. *New England Journal of Medicine*, 367(7):585–590, 2012.

[115] V Oduguwa and R Roy. Bi-level optimisation using genetic algorithm. In *Artificial Intelligence Systems, 2002.(ICAIS 2002). 2002 IEEE International Conference on*, pages 322–327. IEEE, 2002.

[116] Takayuki Osogami. Robustness and risk-sensitivity in markov decision processes. In *Advances in Neural Information Processing Systems*, pages 233–241, 2012.

[117] David L Paterson, Susan Swindells, Jeffrey Mohr, Michelle Brester, Emanuel N Vergis, Cheryl Squier, Marilyn M Wagener, and Nina Singh. Adherence to protease inhibitor therapy and outcomes in patients with hiv infection. *Annals of internal medicine*, 133(1):21–30, 2000.

[118] John D Piette, Jeremy B Sussman, Paul N Pfeiffer, Maria J Silveira, Satinder Singh, and Mariel S Lavieri. Maximizing the value of mobile health monitoring by avoiding redundant patient reports: Prediction of depression-related symptoms and adherence problems in automated health assessment services. *Journal of medical Internet research*, 15(7), 2013.

[119] Michael Pignone, Stephanie Earnshaw, Mark J Pletcher, and Jeffrey A Tice. Aspirin for the primary prevention of cardiovascular disease in women: a cost-utility analysis. *Archives of internal medicine*, 167(3):290–295, 2007.

[120] Michael Pignone, Stephanie Earnshaw, Jeffrey A Tice, and Mark J Pletcher. Aspirin, statins, or both drugs for the primary prevention of coronary heart disease events in men: a cost–utility analysis. *Annals of Internal Medicine*, 144(5):326–336, 2006.

[121] José C Pinheiro and Douglas M Bates. *Linear mixed-effects models: basic concepts and examples*. Springer, 2000.

[122] HA Quigley and AT Broman. The number of people with glaucoma worldwide in 2010 and 2020. *British Journal of Ophthalmology*, 90(3):262–267, 2006.

[123] Nilam Ram and Kevin J Grimm. Methods and measures: Growth mixture modeling: A method for identifying differences in longitudinal change among unobserved groups. *International Journal of Behavioral Development*, 33(6):565–576, 2009.

[124] Nader Rifai, G Russell Warnick, and Marek H Dominiczak. *Handbook of lipoprotein testing*. Amer. Assoc. for Clinical Chemistry, 2000.

[125] Eric M Riles, Anuja V Jain, and A Mark Fendrick. Medication adherence and heart failure. *Current cardiology reports*, 16(3):1–6, 2014.

[126] R Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.

[127] R Tyrrell Rockafellar and Stanislav Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7):1443–1471, 2002.

[128] Véronique L Roger, Alan S Go, Donald M Lloyd-Jones, Emelia J Benjamin, Jarett D Berry, William B Borden, Dawn M Bravata, Shifan Dai, Earl S Ford, Caroline S Fox, et al. Heart disease and stroke statistics–2012 update: a report from the american heart association. *Circulation*, 125(1):e2, 2012.

[129] Allison B Rosen, Mary Beth Hamel, Milton C Weinstein, David M Cutler, A Mark Fendrick, and Sandeep Vijan. Cost-effectiveness of full medicare coverage of angiotensin-converting enzyme inhibitors for beneficiaries with diabetes. *Annals of Internal Medicine*, 143(2):89–99, 2005.

[130] Pavel Roshanov, Hertzel Gerstein, Dereck Hunt, Rolf Sebaldt, and R Haynes. Impact of a computerized system for evidence-based diabetes care on completeness of records: a before–after study. *BMC Medical Informatics and Decision Making*, 12(1):63, 2012.

[131] Roberto Rossi, Onur A Kilic, and S Armagan Tarim. Piecewise linear approximations for the static–dynamic uncertainty strategy in stochastic lot-sizing. *Omega*, 50:126–140, 2015.

[132] Rachel Sacks-Davis, Emma McBryde, Jason Grebely, Margaret Hellard, and Peter Vickerman. Many hepatitis c reinfections that spontaneously clear may be undetected: Markov-chain monte carlo analysis of observational study data. *Journal of The Royal Society Interface*, 12(104):20141197, 2015.

[133] Burhaneddin Sandıkçı, Lisa M Maillart, Andrew J Schaefer, Oguzhan Alagoz, and Mark S Roberts. Estimating the patient's price of privacy in liver transplantation. *Operations Research*, 56(6):1393–1410, 2008.

[134] Valérie Santschi, Arnaud Chiolero, Gilles Paradis, April L Colosimo, and Bernard Burnand. Pharmacist interventions to improve cardiovascular disease risk factors in diabetes a systematic review and meta-analysis of randomized controlled trials. *Diabetes care*, 35(12):2706–2717, 2012.

[135] Sergey Sarykalin, Gaia Serraino, and Stan Uryasev. Value-at-risk vs. conditional value-at-risk in risk management and optimization. *Tutorials in Operations Research. INFORMS, Hanover, MD*, 2008.

[136] Andrew J Schaefer, Matthew D Bailey, Steven M Shechter, and Mark S Roberts. Modeling medical treatment using markov decision processes. *Operations Research and Health Care*, pages 593–612, 2005.

[137] Greggory J Schell, Mariel S Lavieri, Jonathan E Helm, Xiang Liu, David C Musch, Mark P Van Oyen, and Joshua D Stein. Using filtered forecasting techniques to determine personalized monitoring schedules for patients with open-angle glaucoma. *Ophthalmology*, 121(8):1539–1546, 2014.

[138] Greggory J Schell, Mariel S Lavieri, Joshua D Stein, and David C Musch. Filtering data from the collaborative initial glaucoma treatment study for improved identification of glaucoma progression. *BMC medical informatics and decision making*, 13(1):137, 2013.

[139] Steven M Shechter, Oguzhan Alagoz, and Mark S Roberts. Irreversible treatment decisions under consideration of the research and development pipeline for new therapies. *IIE Transactions*, 42(9):632–642, 2010.

[140] Steven M Shechter, Matthew D Bailey, and Andrew J Schaefer. A modeling framework for replacing medical therapies. *IIE Transactions*, 40(9):861–869, 2008.

[141] Steven M Shechter, Matthew D Bailey, Andrew J Schaefer, and Mark S Roberts. The optimal time to initiate hiv therapy under ordered health states. *Operations Research*, 56(1):20–33, 2008.

[142] E.K. Silverman, H.A. Chapman, J.M. Drazen, S.T. Weiss, B. Rosner, E.J. Campbell, W.J. O'DONNELL, J.J. Reilly, L. Ginns, S. Mentzer, et al. Genetic epidemiology of severe, early-onset chronic obstructive pulmonary disease. Risk to relatives for airflow obstruction and chronic bronchitis. *American journal of respiratory and critical care medicine*, 157(6):1770, 1998.

[143] Horst D Simon and Shang-Hua Teng. How good is recursive bisection? *SIAM Journal on Scientific Computing*, 18(5):1436–1445, 1997.

[144] Ankur Sinha, Pekka Malo, Anton Frantsev, and Kalyanmoy Deb. Finding optimal strategies in a multi-period multi-leader–follower stackelberg game using an evolutionary algorithm. *Computers & Operations Research*, 41:374–385, 2014.

[145] Peter C Smith. User charges and priority setting in health care: balancing equity and efficiency. *Journal of health economics*, 24(5):1018–1029, 2005.

[146] Peter C Smith. Incorporating financial protection into decision rules for publicly financed healthcare treatments. *Health economics*, 22(2):180–193, 2013.

[147] Kut C So and Christopher S Tang. Modeling the impact of an outcome-oriented reimbursement policy on clinic, patients, and pharmaceutical firms. *Management Science*, 46(7):875–892, 2000.

[148] Justin W Timbie, Rodney A Hayward, and Sandeep Vijan. Variation in the net benefit of aggressive cardiovascular risk factor control across the us population of patients with diabetes mellitus. *Archives of internal medicine*, 170(12):1037–1044, 2010.

[149] Van-Anh Truong and David Yao. Analytical models for designing pharmaceutical contracts. *Working Paper*.

[150] J Uman, DH Au, and AC Melzer. Predictors of decreased medication adherence for hypertensive patients. *Am J Respir Crit Care Med*, 189:A1410, 2014.

[151] Brian W Ward, Jeannine S Schiller, and Richard A Goodman. Peer reviewed: Multiple chronic conditions among us adults: A 2012 update. *Preventing chronic disease*, 11, 2014.

[152] Gerard J Wedig and Ming Tai-Seale. The effect of report cards on consumer choice in the health insurance market. *Journal of Health Economics*, 21(6):1031–1048, 2002.

[153] ZHAO Wei. Penalty function method for bi-level multiobjective programming. *Acta Automatica Sinica*, 24(3):331–337, 1998.

[154] Kennedy Were, Dieu Tien Bui, Øystein B Dick, and Bal Ram Singh. A comparative assessment of support vector regression, artificial neural networks, and random forests for predicting and mapping soil organic carbon stocks across an afromontane landscape. *Ecological Indicators*, 52:394–403, 2015.

[155] JA Whitworth. 2003 world health organization (who)/international society of hypertension (ish) statement on management of hypertension. *Journal of hypertension*, 21(11):1983–1992, 2003.

[156] J Woodcock. The prospects for personalized medicine in drug development and drug therapy. *Clinical Pharmacology & Therapeutics*, 81(2):164–169, 2007.

[157] Reza Yaesoubi and Stephen D Roberts. Payment contracts in a preventive health care system: A perspective from operations management. *Journal of health economics*, 30(6):1188–1196, 2011.

[158] Yafeng Yin. Genetic-algorithms-based approach for bilevel programming models. *Journal of Transportation Engineering*, 126(2):115–120, 2000.

[159] M. Zahari, B.N. Mukesh, J.L. Rait, H.R. Taylor, and C.A. McCarty. Progression of visual field loss in open angle glaucoma in the Melbourne Visual Impairment Project. *Clinical & experimental ophthalmology*, 34(1):20–26, 2006.

[160] Gregory S Zaric, Hui Zhang, and Reza Mahjoub. Modeling risk sharing agreements and patient access schemes. In *Operations Research and Health Care Policy*, pages 295–310. Springer, 2013.

[161] S.L. Zeger, K.Y. Liang, and P.S. Albert. Models for longitudinal data: a generalized estimating equation approach. *Biometrics*, 44(4):1049–1060, 1988.

[162] Hui Zhang, Gregory S Zaric, and Tao Huang. Optimal design of a pharmaceutical price–volume agreement under asymmetric information about expected market size. *Production and Operations Management*, 20(3):334–346, 2011.

[163] Jin Zhang and Qian Zhang. Stackelberg game for utility-based cooperative cognitiveradio networks. In *Proceedings of the tenth ACM international symposium on Mobile ad hoc networking and computing*, pages 23–32. ACM, 2009.

[164] Jingyu Zhang, Brian T Denton, Hari Balasubramanian, Nilay D Shah, and Brant A Inman. Optimization of prostate biopsy referral decisions. *Manufacturing & Service Operations Management*, 14(4):529–547, 2012.

[165] Yue Zhang, Oded Berman, Patrice Marcotte, and Vedat Verter. A bilevel model for preventive healthcare facility network design with congestion. *IIE Transactions*, 42(12):865–880, 2010.

[166] Donna M Zulman, Eve A Kerr, Timothy P Hofer, Michele Heisler, and Brian J Zikmund-Fisher. Patient-provider concordance in the prioritization of health conditions among hypertensive diabetes patients. *Journal of general internal medicine*, 25(5):408–414, 2010.