

Managing Chronic Health Conditions with Limited Resources

by

Luke J. DeRoos

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
2023

Doctoral Committee:

Associate Professor Mariel Lavieri, Chair
Associate Professor David Hutton
Assistant Professor Neehar Parikh
Professor Kerby Shedden
Associate Professor Cong Shi

Luke J. DeRoos

lkbruski@umich.edu

ORCID iD: 0000-0001-7275-7457

© Luke J. DeRoos 2023

ACKNOWLEDGMENTS

I would like to start by acknowledging my advisor, Mariel Lavieri. Your ability to bring out the best in your students is a gift to us all. Thank you for your foresight, passion, and advocacy.

Thank you to my medical collaborators, Neehar Parikh and Joshua Stein, who gave the very precious gift of their time and guided me throughout this dissertation. Your insights turned my abstract ideas into something meaningful. Thanks also to David Zacks for letting me shadow your clinical work, and to Chris Andrews, Jason Miller, and Elliot Tapper for sharing your expertise.

Thank you to David Hutton, Kerby Shedden, and Cong Shi for serving on my dissertation committee. Your guidance on everything from data sources to model development to validation has been invaluable. Thank you to my lab-mates Daniel Otero-Leon, Lauren Czerniak, Gian-Gabriel Garcia, and Wesley Marrero. You have made me a better researcher, communicator, and teacher, and I am honored to count myself among your ranks. Thanks to Julia Coxen, Erkin Otles, Sajjad Seyedsalehi, and the rest of my PhD colleagues. It is a luxury to have such brilliant, compassionate, and funny friends as coworkers. Thank you to the staff in the IOE department for the countless support services and smiles over the years. Thank you to the Rackham Merit Fellowship and the Seth Bonder Foundation for your financial support and for believing in this work.

Thank you to the many friends outside of my department who have supported me throughout my PhD journey. A special thanks to Vincenzo, Julia, Sarah, and Ariel for serving as equal parts therapist, advisor, and friend. Thank you to my family, who have given me nothing less than their unconditional love and support. Thank you to my parents—James, Lucia, Brian, and Shari—who through lived example have taught me the value of community and hard work. Thank you to Margaret for filling my life with joy and purpose. And thank you to Logan, for everything.

Finally, I want to acknowledge anyone who has ever experienced a chronic disease, and the healthcare workers who serve them. A special thank you to the patients and staff at the Kellogg Eye Center and the United Network for Organ Sharing for making these models meaningful. May we all continue to strive for better, together.

TABLE OF CONTENTS

ACKNOWLEDGMENTS	ii
LIST OF FIGURES	vi
LIST OF TABLES	ix
LIST OF APPENDICES	x
LIST OF ACRONYMS	xi
ABSTRACT	xii
CHAPTER	
1 Introduction	1
1.1 Motivation	1
1.2 Chapter 2: Finding the Optimal Treatment Interval	3
1.2.1 Key contributions	4
1.3 Chapter 3: Synchronizing the Treatment of Multiple Chronic Conditions	5
1.3.1 Key contributions	5
1.4 Chapter 4: Increasing Organ Donation Rates	6
1.4.1 Key contributions	7
2 Finding the Optimal Treatment Interval	8
2.1 Introduction	8
2.1.1 Treat-and-extend	9
2.1.2 Chapter outline	10
2.2 Relevant Literature	10
2.2.1 Markov Decision Processes	10
2.2.2 Chronic disease treatment scheduling	11
2.2.3 Age-related macular degeneration	12
2.3 Modeling Approach	12
2.3.1 Model notation	13
2.3.2 State transitions	14
2.3.3 Current beliefs	15
2.3.4 Optimality equation	15
2.4 Analytical Results	16

2.4.1	Calculating the optimal policy	17
2.4.2	Identifying the MSTI	21
2.4.3	Relaxing the stationary MSTI assumption	25
2.4.4	Incorporating a patient health state	27
2.5	Numerical Results	30
2.5.1	Data sources	30
2.5.2	Model parameterization	30
2.5.3	Example patient	32
2.5.4	Optimal policies for population-level initial beliefs	34
2.5.5	Population level analysis	35
2.5.6	Population level results	36
2.6	Discussion	39
3	Synchronizing the Treatment of Multiple Chronic Conditions	41
3.1	Introduction	41
3.1.1	Chapter outline	42
3.2	Relevant Literature	42
3.3	Modeling Approach	43
3.3.1	Notation	44
3.3.2	State space and state transitions	45
3.3.3	Optimality equation	45
3.4	Analytical Results	45
3.4.1	Assumptions	46
3.4.2	Synchronizing treatment	46
3.4.3	Coordinating two conditions	51
3.4.4	Using a heuristic for large state spaces	58
3.5	Numerical Results	62
3.5.1	Example patient	63
3.5.2	Heuristic and baseline policy analysis	63
3.5.3	Synchronization threshold analysis	66
3.5.4	Population level analysis	70
3.6	Discussion	71
4	Increasing Organ Donation Rates	74
4.1	Introduction	74
4.2	Increasing Donation via Presumed Consent	75
4.2.1	Modeling approach	75
4.2.2	Numerical results	77
4.2.3	Discussion	82
4.3	Increasing Donation via “Ineligible” Donors	85
4.3.1	Modeling approach	86
4.3.2	Numerical results	87
4.3.3	Discussion	98
4.4	Discussion	102
5	Conclusion	104

5.1 Future Work	105
5.2 Concluding Remarks	107
APPENDICES	108
BIBLIOGRAPHY	122

LIST OF FIGURES

FIGURE

1.1	An overview of the dissertation structure.	2
1.2	An example of a patient with two chronic conditions with MSTIs of 3 and 4 periods, respectively. Even with known maximum safe treatment intervals, optimizing the treatment of multiple chronic conditions is difficult. Patients could save a visit by synchronizing treatment, but this synchronization would require additional treatment costs long-term.	6
2.1	The sequence of events during a decision epoch in the MSTI exploration model. . . .	13
2.2	A description of how the lower and upper state indices l and u are updated after observing the outcome of the previous treatment interval.	15
2.3	A visual representation of the proof for Theorem 2.4.1. We use Lemma 2.4.1 to show how to calculate the cost-to-go for state spaces with a single available action. We then show that the cost-to-go of any action can be calculated as a function of these single-state cost-to-go values $V(P, i, i)$	17
2.4	The proportion of Kellogg patients with a given maximum safe treatment interval. . .	31
2.5	An example of a patient following the ordinal MDP using regression-based initial beliefs.	33
2.6	The optimal actions to take under the prevalence-based ordinal MDP and the treat-and-extend policies. As the initial beliefs P are shared across all patients, we show the optimal action in terms of the remaining state parameters l and u . Highlighted cells show where the policies differ.	35
2.7	A visual representation of the discrete event simulation. To estimate the variance in the optimal interval distribution, the simulation was performed on 100 different 50/50 training/testing splits.	36
2.8	Plots comparing the ordinal MDP and treat-and-extend in terms of 1) the average number of weeks with fluid and 2) the average number of visits until reaching a stationary interval. Results are shown across r/c ratios and initial prediction methods. AUC - area under the curve	37
2.9	Plots comparing the ordinal MDP and treat-and-extend in terms of 1) the average number of weeks with fluid and 2) the average number of visits until reaching a stationary interval. Results are shown across discount factors and initial prediction methods. AUC - area under the curve	38
3.1	The sequence of events in each decision epoch.	44

3.2	An example treatment schedule following π_1 for a patient with $M_1 = 3$ and $M_2 = 11$. Assuming that the patient receives treatment for both conditions at time $t = 0$, this specifically represents the treatment cycle corresponding to the cost-to-go value Φ^{π_1} defined in Appendix B.	53
3.3	An overview of the proof of Theorem 3.4.2. Using the policy iteration method of solving dynamic programs, we select π_1 as our initial policy. We then check whether we can improve upon π_1 , which depends on the problem parameters. The conditions under which we cannot improve upon π_1 are the conditions under which π_1 is the optimal policy.	53
3.4	Plots comparing the performance of the baseline decision policies versus the optimal policy. Performance is measured in terms of regret (i.e. the percent increase in cost compared to the optimal decision policy).	65
3.5	Plots comparing the performance of the baseline decision policies versus the optimal policy. Performance is measured in terms of regret (i.e. the percent increase in cost compared to the optimal decision policy).	67
3.6	The full (\bar{c}_0) and joint-start ($\bar{\bar{c}}_0$) synchronization thresholds across a range of maximum safe treatment intervals for an AMD patient receiving bevacizumab.	68
3.7	The full (\bar{c}_0) and joint-start ($\bar{\bar{c}}_0$) synchronization thresholds for an AMD patient with MSTIs of 4 and 7 weeks. For each plot, we assume baseline treatment costs of $c_1 = c_2 = \$50$, and a discount factor of $\delta = 0$, and vary only the parameter of interest.	69
3.8	The distribution of MSTI combinations for patients in the Kellogg dataset. Blank spaces indicate that no patients had that MSTI combination.	70
3.9	The discounted lifetime costs of following each baseline policy across medications.	72
4.1	An overview of the waitlist model dynamics.	76
4.2	Reduction in the number of removals from the waitlist due to death or illness across all organs from 2004-2014 with a 5% increase in donors with presumed consent, stratified by allocation.	78
4.3	Removals from the waitlist due to death or illness from 2004-2014 with a 5% increase in donors with presumed consent, stratified by organ and allocation.	80
4.4	Reduction in the number of waitlist removals due to death or illness by organ	81
4.5	Overall number of waitlist candidates stratified by allocation.	81
4.6	Histograms of ineligible donor use rates across OPOs, stratified by organ.	89
4.7	Kaplan Meier curves for transplant graft survival from January 2008 through November 2020, stratified by organ and donor eligibility.	90
4.8	Kaplan Meier curves for transplant patient survival from January 2008 through November 2020, stratified by organ and donor eligibility.	91
4.9	Kaplan Meier curves for post-transplant graft survival of deceased brain death donations, stratified by organ type and donor eligibility.	93
4.10	Kaplan Meier curves for post-transplant graft survival, stratified by organ type and donor death type (DCD: deceased cardiac death, DBD: deceased brain death).	94
4.11	Kaplan Meier curves for post-transplant graft survival of ineligible donations, stratified by organ type and donor death type (DCD: deceased cardiac death, DBD: deceased brain death).	95

4.12	Selected hazard ratios calculated via Cox regression for transplant graft survival from January 2008 through November 2020. Recipient age is scaled to be in decades. Recipient age, ethnicity, and BMI were all significantly associated with graft survival ($p \leq 0.05$). Abbreviations: eth.: ethnicity.	96
4.13	Selected hazard ratios calculated via Cox regression regarding the interaction effect of donor ineligibility with recipient age, ethnicity, sex, and BMI. Results are stratified by organ.	97
4.14	Map showing the estimated annual increase in life-years gained by organ procurement organization (OPO), if all OPOs with ineligible donation use rates below the 75th percentile increased their use to meet the 75th percentile. The increase shown is across all organs, based on annual use from January 2008 through November 2020.	99
C.1	Hazard ratios calculated via Cox regression regarding the association of organ procurement organization (OPO) with graft survival after controlling for donor eligibility, as well as recipient age, ethnicity, sex, and BMI. OPO names have been de-identified using a randomly assigned three-digit code. OPO 120 was randomly selected as the baseline OPO.	121

LIST OF TABLES

TABLE

3.1	The discounted cost of following different decision policies for a joint-start bilateral AMD patient being treated with bevacizumab.	63
4.1	The size of the organ waitlists on January 1st of each respective year.	77
4.2	Mean monthly reduction in removals due to sickness or death by organ, 2004-2014 . . .	79
4.3	Percent change in the number of waitlist candidates from 2004-2014 under a random allocation policy	82
4.4	Average annual estimated life years gained by organ and presumed consent impact for both random and ideal allocations	83
4.5	Organ donor population description. P-values are the result of χ^2 -test (categorical) and t-test (continuous) variable comparisons between eligible and ineligible subgroups. SD – standard deviation.	88
4.6	Estimated annual increases in transplants and life-years gained associated with increasing ineligible donor use under a range of scenarios. To calculate the percentile match results, all OPOs with ineligible donor use rates below the given percentile had their rates increased to match the percentile.	98

LIST OF APPENDICES

A Supporting Material for Chapter 2 108
B Supporting Material for Chapter 3 109
C Supporting Material for Chapter 4 117

LIST OF ACRONYMS

AMD age-related macular degeneration

anti-VEGF anti-vascular endothelial growth factor

BMI body mass index

COD cause of death

DCD donation after cardiac death

DBD donation after brain death

HR hazard ratio

MDP Markov Decision Process

MELD model of end-stage liver disease

MSTI maximum safe treatment interval

OCT optical coherence tomography

OPO organ procurement organization

OPTN Organ Procurement and Transplantation Network

SES socioeconomic status

STAR Standard Transplant and Research

UNOS United Network of Organ Sharing

ABSTRACT

In this work, we aim to help patients and providers manage chronic health conditions using operations research. Providing quality care to patients means managing limited resources, including time, money, and medication. We explore two approaches to minimize the burden of resource scarcity on patient outcomes: 1) reducing resource demand and 2) increasing resource availability. First, we focus on reducing resource demand by optimizing the treatment regimen required to prevent disease progression. We provide a means of finding the optimal treatment interval for a patient in order to minimize the number of treatments and clinic visits without compromising patient health. We then provide a framework for optimizing the treatment of multiple chronic conditions and describe when it is optimal to synchronize treatment across conditions. We highlight the usefulness of these treatment planning models using a case study on patients with neovascular age-related macular degeneration, a chronic eye disease. Second, we seek to maximize the availability of chronic disease treatment in situations where patients face resource scarcity—such as in the case of organ transplantation. For patients on the transplant waiting list, the need for transplants far outpaces donation rates. We describe two policies designed to improve patient outcomes by increasing donation rates. We believe this work offers both theoretical and practical value to the field of healthcare operations research. By making chronic disease treatment more efficient and more readily available, we hope that the decision policies presented here drive meaningful change and make a lasting difference in the lives of patients.

CHAPTER 1

Introduction

In the United States, 6 in 10 adults have at least one chronic health condition, and 4 in 10 have two or more. [16] The Centers for Disease Control broadly define a chronic disease as a condition that lasts 1 year or more, and requires ongoing medical attention or limits the activities of daily living, or both. [16] On a population level, managing chronic conditions accounts for 90% of healthcare spend in the United States. [16] On an individual patient level, a chronic disease diagnosis can mean years of frequent, expensive visits to a healthcare provider. In all cases, both patients and providers are faced with the challenge of improving or maintaining patient health with a limited set of resources.

The resources required to manage a chronic condition come in many forms, such as time, money, and medication. In the face of resource scarcity, two general approaches of improving outcomes are 1) reducing resource demand or 2) increasing resource availability. In this work, we seek to improve patient outcomes using both of these approaches. First, we develop a Markov Decision Process (MDP) to identify the optimal treatment interval for a given condition—an interval which minimizes the number of required treatments without compromising on patient health. By minimizing the number of visits, we also minimize the resources required to complete these visits. Second, we develop a dynamic programming model to optimize treatment timing for patients with multiple chronic conditions. We demonstrate when and how to synchronize treatment in order to further reduce visit costs. Third, we discuss methods to increase resource availability for an important chronic disease treatment modality: organ transplant. Figure 1.1 provides a visual representation of this structure.

1.1 Motivation

While our methods can be applied to chronic conditions more broadly, in this work we focus on two specific chronic conditions to provide context: neovascular age-related macular degeneration (AMD) and organ failure.

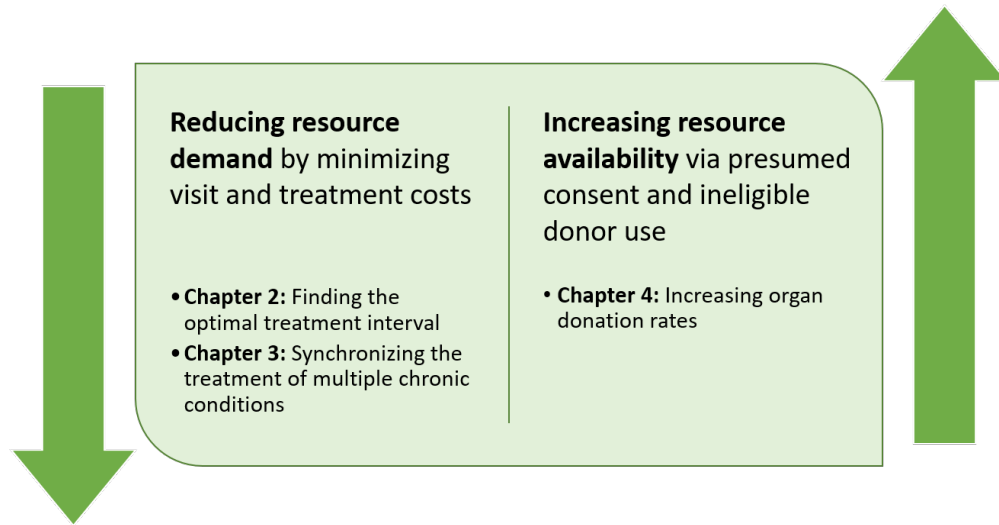


Figure 1.1: An overview of the dissertation structure.

Neovascular AMD is a chronic eye disease that places an intense burden on diagnosed patients. AMD is the most common cause of blindness in people over 55 years old, and nearly 170,000 adults in the United States are diagnosed with neovascular AMD each year [13]. Neovascular AMD is characterized by the growth of excess blood vessels in the eye, which leak fluid and distort the surface of the retina. This leads to blurred or distorted vision, which reduces one’s ability to drive or read, and significantly affects overall quality of life. To assess the progression of AMD and identify retinal distortion, clinicians rely on optical coherence tomography (OCT) scans, which provide quantitative measurements of retinal thickness. Clinicians primarily treat AMD through the injection of anti-vascular endothelial growth factor (anti-VEGF) medication directly into an affected eye [98]. To prevent fluid from recurring, patients may require injections as frequently as once per month, imposing a substantial burden on patients, especially patients with significant vision loss who may not be able to drive themselves to appointments. However the fluid recurrence rate and the response to medication differs significantly among patients and even between eyes. As a result, it is challenging for clinicians to determine the optimal timing of treatment for a given patient. Administering injections too frequently imposes unnecessary discomfort and increases the risk of serious bleeding, infection, and cataracts—all without providing additional benefit to the patient’s long-term vision. In contrast, administering injections too infrequently puts a patient at risk of fluid buildup in the eye, resulting in persistent retinal damage and loss of vision. In this work, we provide an optimization-based framework to address these scheduling challenges. Our hope is to provide both clinicians and patients with effective and interpretable tools to help minimize the burden of this demanding disease.

Organ transplantation is a life-saving and cost-effective intervention for patients with organ

failure. [6] However, the supply of organs in the United States cannot meet the need of patients on the waitlist, resulting in prolonged morbidity and increased mortality of waitlisted patients. [48] In the US, over 7,000 waitlisted candidates die each year before ever receiving a transplant. Broadly, there are two ways to increase the number of available donors: 1) optimizing donor yield and 2) increasing the donor pool. Optimizing donor yield requires the use of best practices and emerging technologies (i.e. machine perfusion) in donor consent and organ retrieval. [96] Even with these mechanisms, the donor pool is limited and will likely shrink in coming years due to changing US demographics and population health (e.g. increasing obesity, ageing population), further exacerbating the disparity between donors and recipients. [70] Additionally, technological advancements often take years of extensive validation before they can be safely and effectively implemented. In contrast, an immediate method of reducing the gap between organ donation and demand is to increase donation rates. In this work, we discuss two methods of increasing organ donation and characterize the potential impact of these policies on transplant candidate outcomes. Our goal is to help guide the efforts of policymakers by providing data-driven estimates of policy outcomes on the many patients in desperate need of an organ transplant.

1.2 Chapter 2: Finding the Optimal Treatment Interval

The goal of Chapter 2 is to minimize the cost of treatment for patients with a chronic disease, without compromising on their long-term health. We focus on a class of chronic conditions (which includes AMD) that have what we refer to as a maximum safe treatment interval (MSTI). The defining characteristic of conditions with an MSTI is intuitive: as long as a condition is treated within some interval (the MSTI), then we do not expect disease progression. For AMD, patients receive regular injections to prevent neovascular fluid. As long as they receive injections within a given interval, fluid is prevented. If they go too long between injections, then fluid reoccurs and the patient can have long-term vision damage. If the MSTI of a condition is known then, for a single condition, the optimal treatment schedule is to treat exactly at the MSTI. This minimizes the number of visits without disease progression. However, the MSTI of a condition varies across patients and conditions, and is initially unknown. Identifying the MSTI can be challenging, and the searching for it can expose patients to unnecessary costs or health risks if not managed appropriately.

In our work, we provide an optimization approach to finding the MSTI of a chronic condition. We model this problem as an infinite horizon MDP, designed to balance the exploration/exploitation trade-off found in the MSTI search.

1.2.1 Key contributions

- **Ordinal action space:** In the traditional treatment scheduling literature, scheduling decisions are based on *whether* to treat in each time period. In our model, we instead decide *when* to return for the following treatment. Our decision points are each clinical visit (as opposed to each time period), and our action space is the number of time periods between now and the next injection. This paradigm shift gives our action space the key property of ordinality. The primary benefit of ordinality among actions is that we are able to learn about the expected outcome of *multiple* actions from a *single* decision.
- **Offline solution:** By leveraging the ordinality of our action space, we are able to derive closed-form solutions for the optimal action in each decision period. This has several benefits. First, it reduces computational complexity, which is a ubiquitous problem when solving MDPs. Second, it provides each action with an “index” score, representing the value of picking that action. These index scores improve interpretability from a clinical perspective, and provide clinicians with a menu of treatment options. For example, if the optimal recommended action is infeasible for some reason (e.g. the patient has a scheduling constraint) the clinician can look at the other options and see what feasible actions might have minimal loss of optimality. Third, the index policy allows us to determine under which conditions commonly used MSTI exploration policies are optimal policies. Fourth, the index policy allows us to generate interpretable and generalizable policies for clinical scenarios in which the clinician does not have any prior knowledge about a patient’s health state.
- **Data robust policies:** By leveraging our index value approach, we are able to provide population level policies for situations in which clinicians have no previous knowledge about the patient. Even without personalized patient data, we are able to improve upon current clinical protocols with simple recommendations, such as increasing the treatment interval by 1-week intervals when exploring (versus the current practice of 2-week extensions).
- **Long-term treatment guidelines:** In general, a patient’s MSTI for a condition like AMD is stationary. [57, 39, 22] Consequently, once patients reach a MSTI treatment interval, they are treated at this interval for the remainder of care. However, to generalize our framework, we consider the case in which a patient’s MSTI changes over time. We provide a mathematical approach for when clinicians should test longer treatment intervals to determine if the patient could receive added benefit from an interval extension.
- **Practical clinical benefit:** As a case study, we apply our framework to the treatment of AMD using patient data from the Kellogg Eye Center. We demonstrate that following our ordinal MDP could reduce patient exposure to symptoms by up to 36% while finding the MSTI

7% faster than the current standard of care. Additionally, several new treatment modalities for AMD are anticipated to extend the maximum treatment interval from 12 weeks to one on the order of months or years. [45] For such extended time periods, there are no recommendations for finding the MSTI, and current protocols will be too slow to be practically useful. The ordinal MDP can be used to inform the appropriate exploration policy as the range of available intervals grows.

1.3 Chapter 3: Synchronizing the Treatment of Multiple Chronic Conditions

Our goal in Chapter 3 is to build upon Chapter 2 and minimize treatment costs for patients with *multiple* chronic conditions. For a single condition, once the MSTI is identified, the optimal treatment schedule is straightforward. However, for patients with multiple chronic conditions, knowing the MSTI is often not enough to provide clear scheduling recommendations. Consider the example shown in Figure 1.2. Here, a patient has two chronic conditions with MSTIs of 3 and 4 periods, respectively. To minimize treatment costs, the clinician could schedule treatment at the MSTIs. However, simply by scheduling the second condition 1 period sooner, the clinician could synchronize the patient's treatment and save them a visit to the clinic, which often means saving the patient significant time and money. However, this would require additional condition 2 treatment costs long-term. In sum, clinicians are constantly managing a trade-off between minimizing visit costs and minimizing treatment costs.

In this work, we model the multiple condition scheduling problem as a dynamic program to determine optimal scheduling policies. Beyond finding the optimal scheduling policy, we use our model to answer a key practical question: when is it worth synchronizing the treatment of multiple chronic conditions?

1.3.1 Key contributions

- **Multiple condition model:** On the most basic level, this work is one of the first applications of treatment optimization for *multiple* conditions. Managing multiple chronic conditions invokes several important research questions. First, it is unclear how much utility is lost by scheduling and diagnosing each condition independently instead of jointly. Our model helps us understand when it is beneficial to schedule treatment for multiple conditions in the same visit. In practice, scheduling decisions are almost always made considering all conditions, as they each have a major impact on patient quality of life. This model better aligns mathematical decision making to the everyday decisions made by clinicians and patients everywhere.

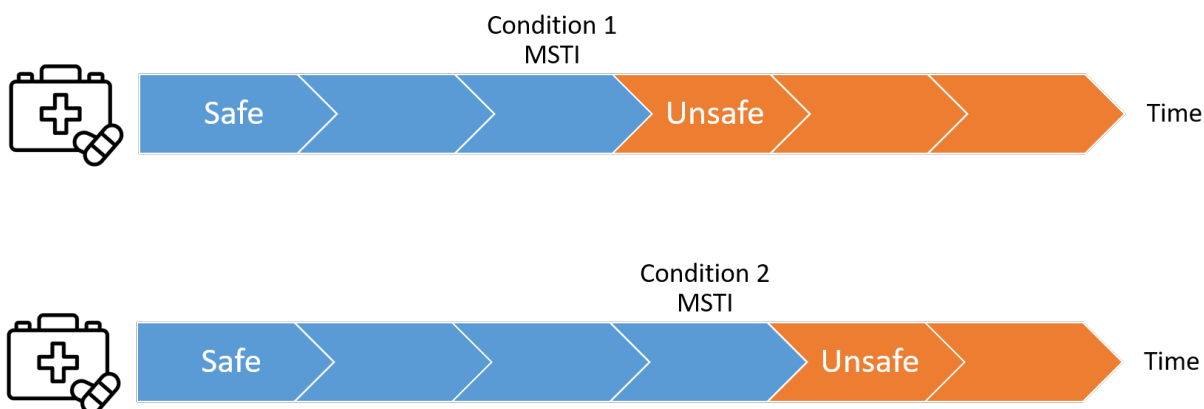


Figure 1.2: An example of a patient with two chronic conditions with MSTIs of 3 and 4 periods, respectively. Even with known maximum safe treatment intervals, optimizing the treatment of multiple chronic conditions is difficult. Patients could save a visit by synchronizing treatment, but this synchronization would require additional treatment costs long-term.

- **Synchronization guidelines:** We analytically define conditions under which it is optimal to synchronize treatment. We show that if a patient’s visit cost is high enough, it is always optimal to synchronize treatment. Additionally, we provide the optimal synchronization policy in this case. For patients with 2 chronic conditions, we provide a closed form solution for the visit cost threshold above which the optimal policy is to synchronize treatment.
- **Robust heuristic:** We provide an easily solvable heuristic for synchronizing treatment in the 2-condition case, and analytically bound the potential regret of following this heuristic. We numerically demonstrate that the heuristic is almost always optimal for patients diagnosed with advanced AMD.
- **Practical clinical benefit:** Using data from the Kellogg Eye Center, we demonstrate that our optimization approach could save patients an average of \$1,189 in direct medical costs. Extrapolating this to a national level, this could result in approximately \$582 million in savings. In addition to the direct medical cost savings, a reduced number of visits can also have a dramatic impact on patient quality of life.

1.4 Chapter 4: Increasing Organ Donation Rates

Our goal in Chapter 4 is to increase the availability of treatment for patients awaiting an organ transplant. Where Chapters 2 and 3 focus on reducing the treatment demand for chronic dis-

ease patients, this chapter instead focuses on increasing the supply of a chronic disease treatment modality. We do this by developing statistical models to assess the impact of policies designed to increase organ donation rates in the United States. Implementing policies to close the donation gap would offer a major societal benefit. However, fully implementing any policy is challenging, especially those on a population level. An understanding of how a policy would impact the transplant candidate population can help focus the efforts of policymakers and improve the likelihood of implementation.

1.4.1 Key contributions

- **Presumed consent simulation:** We build a simulation model estimating the impact of a presumed consent policy on US donation rates. Under this policy, which has been implemented in several countries, willingness to donate is the default option. [77] We quantify the potential effect of presumed consent on waitlist mortality and size, and show that implementing this policy could provide as many as 34,000 life-years annually.
- **Ineligible donor survival analysis:** We analyze the use of so-called “ineligible” or non-standard organ donors in the United States. We show that several transplant organizations successfully use ineligible donors with no or minimal detriment to recipient survival. We quantify the potential benefit of standardizing ineligible organ donor use nationally, showing a potential increase of 16,000 life-years annually.
- **Ineligible donor use model:** We build a beta regression model to understand contextual factors associated with ineligible donor use. This helps us better understand the heterogeneity in donor use, and provides a starting point for best practice sharing among transplant organizations.

Decision making in healthcare is incredibly complex, and each decision can have a major effect on a patient’s quality of life. The work presented here is a stepping stone in a long line of research on chronic disease management. We hope that this work has a meaningful impact on the way we approach treatment planning, and ultimately on patient lives.

CHAPTER 2

Finding the Optimal Treatment Interval

2.1 Introduction

For many patients, a chronic disease diagnosis means years of frequent visits to healthcare professionals to receive expensive, yet essential treatment. For clinicians, appropriately timing treatment is crucial to providing quality care. Undertreating a condition can lead to poor health outcomes for patients and increase long-term treatment needs, while overtreating can result in unnecessary side effects, risks, and increased costs. Many conditions, such as AMD, have what we refer to as a “maximum safe treatment interval” or MSTI. Finding the MSTI is important, because once discovered, scheduling exactly at the MSTI is optimal—it minimizes the treatment frequency without compromising the patient’s long-term health. However, even within the same condition, it is challenging for clinicians to find the MSTI as it can vary significantly among patients [23]. In this chapter we introduce the ordinal MDP framework as a solution to this challenge. We demonstrate how the ordinal MDP can quickly identify the appropriate treatment schedule for patients while reducing exposure to unnecessary risk.

The ordinal MDP derives its name from the ordinality of its action space. The action space represents different treatment intervals, which can be sorted by length of time. Traditionally, MDP scheduling models guide decisions around *whether* or not to treat a patient in each decision epoch. In contrast, the ordinal MDP suggests *when* to next treat the patient. By reframing the decision paradigm, we better align the model with the scheduling choices that clinicians face in practice. Additionally, the ordinal action space allows clinicians to learn about the expected outcomes of multiple treatment intervals with only a single scheduling decision, even intervals which the clinician has yet to explore. For example, suppose a clinician schedules treatment three months into the future and, after observing the patient upon return, the clinician learns that a three month interval was too long between treatments. Via the same decision, the clinician has learned that any interval *longer* than three months is also too long between treatments. More generally, this ordinal structure allows individual observations to update the expected rewards of multiple actions in a single

period. This structure allows the ordinal MDP to overcome common challenges such as computational tractability and practical feasibility. In this work, we demonstrate that the ordinal structure allows for a closed form, index-based solution of the optimal decision policy under uncertainty.

We focus on chronic health conditions that have an MSTI, such as AMD, diabetic macular edema, urticaria, and osteoarthritis. [98, 71, 90, 74] For this class of conditions, studies show that while the MSTI varies significantly among patients, it remains consistent for a given patient over the course of treatment. [57, 39, 22] As a result, clinical experts emphasize the importance of identifying personalized fixed-interval treatment schedules when managing this class of conditions. [23, 51] Consistency in individual patient scheduling is important for several reasons. High variability in treatment timing can result in symptom flare-ups between visits and require excess treatment to manage both the short- and long-term effects of these symptoms. Highly variable treatment intervals also make it difficult for patients and caregivers to manage their schedules, imposing additional financial and psychological burden on people already facing intense physiological obstacles. From a population perspective, this inconsistency also makes it difficult to identify and share best practices among providers. Currently, clinicians rely on heuristics to find a patient’s MSTI. [98, 71, 90, 74] With this work, we instead provide an optimization framework for safely and efficiently finding the optimal treatment interval.

2.1.1 Treat-and-extend

In current practice, a majority of clinicians use the “treat-and-extend” protocol to manage the treatment of neovascular AMD. [23, 98] Treat-and-extend searches for a patient’s MSTI by first scheduling patient treatment using the shortest recommended time between injections (4 weeks). This initial 4-week interval is based on the efficacy of the drug as observed in clinical trials. [98] Immediately prior to each treatment, the clinician also observes whether or not retinal fluid is present. Once the patient reaches a baseline response level in which fluid is consistently not present, the clinician begins extending the treatment interval. The clinician first extends the interval to 6 weeks. If the patient returns with no symptoms, the clinician continues to extend this interval by 2-week increments until symptoms arise or until the patient reaches the maximum recommended interval of 12 weeks. [98] Then, the clinician reverts to the longest tested interval in which the patient had no symptoms. If the patient did not reach the 12-week maximum, the clinician then performs a finer search by extending the interval using 1-week increments until symptoms again arise or until the patient would reach an interval known to cause symptoms. At this point, the search is complete, and the clinician indefinitely treats the patient using the longest interval which does not cause symptoms.

Healthcare providers use protocols similar to treat-and-extend to manage other chronic con-

ditions where response to medication is stable for a given patient, but varies among individuals. [71, 90, 74] However, these protocols are a “one-size-fits-all” approach, and they do not take advantage of personalized patient information which can improve the speed and safety of exploring available treatment intervals.

2.1.2 Chapter outline

The remainder of this chapter is organized as follows: in Section 2.2 we discuss the literature on scheduling MDPs with a particular focus on their application to chronic disease. We also provide background on neovascular AMD and current treatment methods. In Section 2.3 we outline the formulation of the ordinal MDP. In Section 2.4 we discuss the derivation of the optimal policy under uncertainty, as well as analytical results and insights around finding a patient’s MSTI with varying model parameters. We also discuss a method for determining when to check if a patient’s MSTI has gotten longer over the course of treatment. In Section 2.5 we offer numerical results via a case study on neovascular AMD. Finally, in Section 2.6 we conclude and discuss future research avenues.

2.2 Relevant Literature

The use of operations research to address complex healthcare challenges continues to grow. [20] Even within healthcare, this body of work is broad in both application and approach. To provide appropriate context for this study, we survey the following themes within operations research: 1) Markov Decision Processes, 2) chronic disease treatment scheduling, and 3) AMD management.

2.2.1 Markov Decision Processes

Methodologically, this research builds upon the current MDP literature. Our problem is an example of an exploration-exploitation trade-off. This trade-off is one where a decision maker seeks to learn about available options while still exploiting current knowledge to maximize long-term rewards. Multi-armed bandit (MAB) models are a specific type of MDP that highly useful for managing the exploration-exploitation trade-off. Gittins (1979) established the primary results of the MAB problem in his seminal work, where he showed that under certain conditions, the complex exploration-exploitation problem could be simplified to an index-based policy with a closed form solution. [26] However, these conditions are uncommon in practice, and the literature has expanded to consider many variants of the MAB, such as the restless bandit [95] or the contextual bandit [97]. Many implementations of these alternative MAB formulations do not have computationally tractable solutions, and rely on a vast literature of reinforcement learning heuristics,

such as the ϵ -greedy [94], interval estimation [43], Boltzmann exploration [54], upper confidence bound [5], and Thompson sampling [2] algorithms. Some examples of MABs in healthcare focus on applications such as liver cancer, hepatitis, and multiple sclerosis. [3, 65, 52, 8] Again, these approaches consider whether or how aggressively to treat a patient in each period, and do not incorporate an ordinal action space. In this work, we augment the traditional MAB problem by structuring our model to have dependent arms. Often, independence between arms is considered a core tenet of a MAB, and relaxing this assumption also relaxes the problem to be a general MDP. As a result, we refer to our model as an MDP, despite its close connection to the MAB structure. In this work, we build upon the literature by describing conditions under which the optimal exploration policy is an index policy, even with dependent arms.

2.2.2 Chronic disease treatment scheduling

Operations researchers have studied treatment scheduling across a number of chronic conditions. Perhaps closest to our work are Helm et. al (2015) and Schell et. al (2014), who provide heuristic approaches to calculate the time between visits for patients with glaucoma. [38, 84] However, these approaches exclusively consider the next visit, and do not consider the impact of timing on future visits and long-term patient outcomes. Alagoz et. al (2004), Schaefer et. al (2005), and Denton (2018) provide an overview of how MDPs are typically used in treatment planning, and how they incorporate long-term rewards to generate more holistic scheduling policies. Some specific examples of MDPs in treatment scheduling include Schechter et. al (2008), who study the optimal time to initiate HIV therapy, Ayer et. al (2012), who leverage personalized data to improve the timing of breast cancer screening, and Kamalzadeh et. al (2021), who use a partially observable MDP to optimize the timing of diabetes screening. [87, 7, 44] Researchers have also utilized MDPs to address challenges in organ transplant, head and neck cancer, and chronic kidney disease. [83, 9, 66, 89]

In our framework, we leverage the ordinal relationship between decisions to gain both managerial insights and computational advantages. This ordinal relationship is one that clinicians often anecdotally consider in scheduling decisions, but it is not incorporated in personalized scheduling models. [98] Our paper formally incorporates this clinical use of ordinality into a mathematically rigorous optimization model. Beyond providing significant practical benefits to patients, research also suggests that limiting the scope of scheduling policies to those with stationary intervals long-term has a marginal effect on patient outcomes. [18]

2.2.3 Age-related macular degeneration

The current AMD treatment landscape is not without scheduling policy recommendations. Namely, research has shown the treat-and-extend policy to be a cost-effective means of reducing patient treatment burden without significantly harming health outcomes. [32, 23, 10] In fact, the structure of treat-and-extend was a significant inspiration for the structure of the ordinal bandit. However, treat-and-extend was not developed using an optimization approach, nor does it fully leverage personalized patient information. To the best of our knowledge, there are no studies on whether or not the current policy parameters (e.g. the rate of interval extension) are optimal, or whether current patient health information can be used to improve scheduling outcomes. The ordinal bandit enhances the treat-and-extend policy by answering these questions via an optimization framework.

More generally, machine learning as a means of enhancing AMD treatment has become increasingly popular in recent years. Modern AMD decision making revolves heavily around the interpretation of OCT images, which provide a visual and statistical representation of a patient’s retina, including size, shape, and the presence of abnormalities. Bogunovic et. al (2017) and Schlegl et. al (2018) are among a growing number of works automating the identification of AMD symptoms using machine learning image processing techniques and estimating individual patient response to treatment. [11, 85] Other works, such as Schmidt-Erfurth et. al (2018) and Rohm et. al (2018), have extended this research to estimate long term visual acuity for patients experiencing AMD. [86, 78] While these models are very useful in understanding current patient biomarkers, they do not incorporate the clinical decision making process. Our model expands the AMD literature by offering data-driven recommendations on the timing of treatment.

2.3 Modeling Approach

We model this problem as a discrete time, infinite horizon MDP. The events within each decision epoch are outlined in Figure 2.1. The decision maker seeks to minimize both the time a patient experiences symptoms and the required number of treatments. Ultimately, the decision maker is searching for a patient-specific interval a_M , which represents the longest interval between treatments during which the patient will not experience adverse health effects (i.e. the patient’s MSTI).

In an initial treatment phase, the clinician gathers patient information and generates a set of initial beliefs P , representing her initial beliefs of the true value of a_M . Thereafter, in each visit, the clinician first observes whether or not the patient has experienced symptoms since the last treatment. The clinician uses this observation to update the patient state X_t , which contains information on the efficacy of scheduling intervals observed prior to visit t . The clinician then updates her set

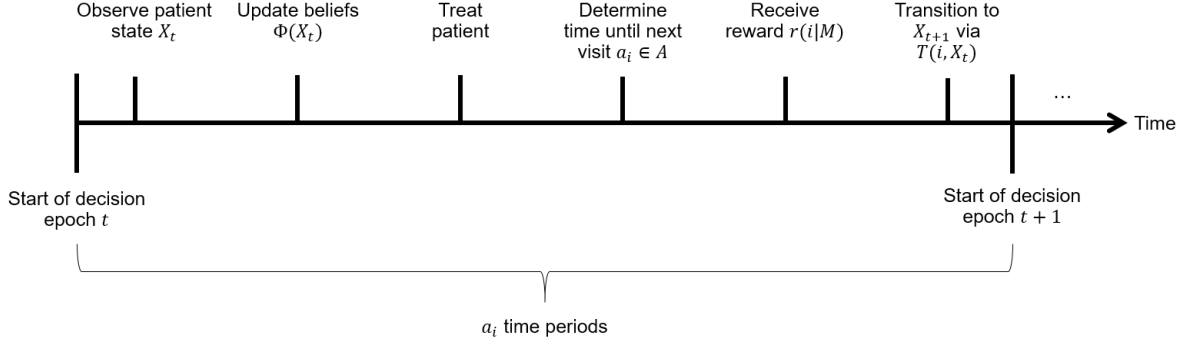


Figure 2.1: The sequence of events during a decision epoch in the MSTI exploration model.

of current beliefs $\Phi(X_t)$, representing her new estimates of the true value of a_M for the individual patient. The clinician then treats the patient. We assume the amount and type of treatment is given according to current clinical guidelines. Given her current beliefs, the clinician then selects an action $a_i \in A$, representing the amount of time until the patient should return for the next treatment. Based on this decision, the patient receives an immediate reward of $r(i|M)$. This reward represents the amount of time the patient spends outside of the clinic without experiencing symptoms. The patient then transitions to a new state X_{t+1} according to a transition matrix $T(i, X_t)$. After a_i time periods, the patient returns for the next treatment, and the process repeats.

With this modeling approach, we specifically consider the class of chronic conditions in which patients have different treatment needs and, after some exploration, eventually follow a personalized, fixed-interval treatment plan. In our case, we assume that each patient has an MSTI. We also assume that the clinician cannot observe the patient between visits. This is appropriate for conditions such as AMD, where it is difficult for the patient to assess the presence of symptoms at home. As chronic conditions often affect patients for the remainder of their lifetime, the decision horizon is generally much longer than individual treatment intervals. We therefore model this as an infinite horizon problem, as is common in the medical decision making literature. [4, 87, 91]

2.3.1 Model notation

We can formalize our model using the following notation:

- t : the decision epochs, $t \in \{1, \dots, \infty\}$.
- A : the set of available actions, $A = \{a_1, \dots, a_N\}$. Here, $N < \infty$ and each $a_i \in A$ is a positive value representing a potential treatment interval length. Without loss of generality, let $i < j \implies a_i < a_j$, where $i, j \in \{1, \dots, N\}$.
- a_M : the longest interval that does not cause adverse health effects for the patient. We also

refer to a_M as the patient's maximum safe treatment interval, or MSTI. Here, $a_M \in A$ is initially unknown to the clinician.

- p_i : the initial belief that action a_i is the MSTI, $p_i := P(a_i = a_M)$
- P : the clinician's set of initial beliefs, $P := \{p_i : a_i \in A\}$.
- l : an index such that $a_l \in A$ is the shortest action that is potentially the MSTI. An alternative definition of l is that it is the index of longest interval $a_l \in A$ that is known to not cause an adverse health state for the patient.
- u : an index such that $a_u \in A$ is the longest action that is potentially the MSTI.
- X : the current state, $X = (P, l, u)$. As this is an infinite horizon problem, we omit the time notation from the state where unambiguous.
- $T(i, X)$: the transition probability matrix given action a_i and state X .
- $\phi(i|X)$: the current belief that action a_i is the MSTI given state X , $P(a_i = a_M|X)$
- $\Phi(X)$: the clinician's set of current beliefs, $\Phi = \{\phi(i|P, l, u) : a_i \in A\}$.
- $r(i|M)$: the immediate reward for selecting action a_i given the MSTI a_M .
- δ : the per-period discount factor, $\delta \in (0, 1)$. Note that a decision epoch may contain multiple time periods.

2.3.2 State transitions

When a patient returns to a clinic for treatment, we assume that the clinician is able to observe whether or not the previously scheduled interval was too long to prevent adverse health outcomes. For example, clinicians may use image scans or blood draws to assess whether there is evidence of disease progression or whether the patient had symptoms reoccur in between treatments. After taking an action a_i , one of two possible outcomes is observed:

1. the treatment interval was shorter than or equal to the MSTI a_M , or
2. the interval was longer than the MSTI a_M .

In scenario 1, the clinician has learned that the MSTI is at least as long as the previous action. Mathematically, the lower index l can be updated to the index i of the previous action a_i . In

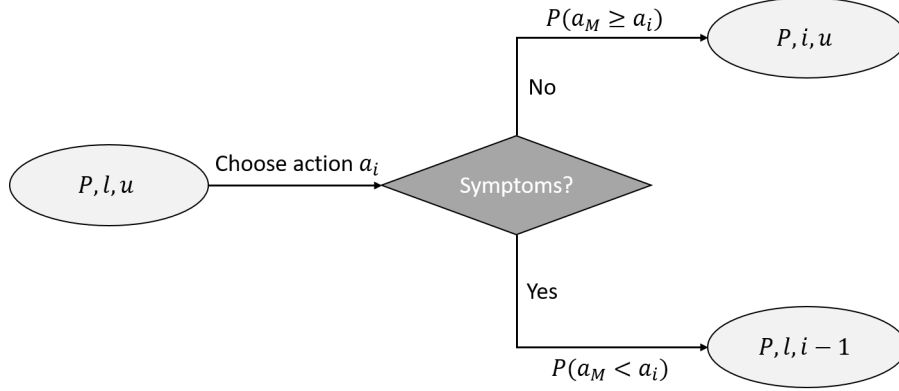


Figure 2.2: A description of how the lower and upper state indices l and u are updated after observing the outcome of the previous treatment interval.

scenario 2 the clinician has learned that the MSTI is strictly less than the previous action. Mathematically, the upper index u can be updated to be $i - 1$. Figure 2.2 is a graphical representation of this transition process.

From Figure 2, we can see that our transition matrix $T(i, X_t)$ is quite sparse: given $X_t = (P, l, u)$ and action a_i , we observe that the subsequent state X_{t+1} can only take on two (l, u) combinations. If no symptoms are observed, then $X_{t+1} = (P, i, u)$, which occurs with probability $P(a_M \geq a_i) = \sum_{j=l}^i \phi(j|P, l, u)$. If symptoms are observed, then $X_{t+1} = (P, l, i - 1)$, which occurs with probability $P(a_M < a_i) = \sum_{j=i+1}^u \phi(j|P, l, u)$.

2.3.3 Current beliefs

The ordinality of A allows us to concisely represent the clinician's current knowledge at any point in time using three parameters (P, l, u) . Using Bayes' rule we have:

$$\phi(i|P, l, u) = \begin{cases} \frac{p_i}{\sum_{i=l}^u p_i} & \text{if } l \leq i \leq u \\ 0 & \text{otherwise} \end{cases} \quad (2.1)$$

From Equation 2.1 we see that all of the information required to calculate our current beliefs is contained in our state, and does not require storing any previous state information.

2.3.4 Optimality equation

The state transitions described in Section 2.3.2 allow for an intuitive formulation of the ordinal bandit optimality equation, written below as Equation 2.2. Because the clinician cannot observe the patient between visits, she is unable to observe exactly when symptoms arise. She is only

aware of the existence of symptoms between treatments based on the patient’s health state upon return. As a result, the immediate rewards associated with changes to the patient’s health state are only partially observable, and the clinician estimates the immediate rewards based on her current beliefs.

While the immediate rewards reflect the expected health state of the patient, the future rewards contain information on both the patient’s future health and the value of knowledge gained from taking a decision. For example, if a clinician selects the shortest action a_l , they know by definition that there will be no adverse health events before the next treatment. However, selecting the shortest action provides no new information about which scheduling interval is the MSTI. As a result, when making this decision, the state space will be the same in the next decision epoch. In contrast, selecting an action that is not the shortest interval *will* provide additional knowledge on the MSTI, regardless of whether or not symptoms occur before the next treatment. Equation 2.2 is written to demonstrate these two scenarios. Additionally, we list the scenarios separately to avoid writing the degenerate case of $V(P, l, l - 1)$, and to provide intuition for the index policy derivation discussed in Section 2.4.1.

$$\begin{aligned}
 V(P, l, u) = & \\
 \max_{d \in \{l, \dots, u\}} & \left\{ \begin{array}{ll}
 \sum_{i=l}^u (r(l|i) \cdot \phi(i|P, l, u)) + \delta^{a_l} V(P, l, u), & \text{if } d = l \\
 \sum_{i=l}^u (r(d|i) \cdot \phi(i|P, l, u)) + \dots \\
 \delta^{a_d} \left(\sum_{i=l}^{d-1} (\phi(i|P, l, u)) \cdot V(P, l, d - 1) + \sum_{i=d}^u (\phi(i|P, l, u)) \cdot V(P, d, u) \right), & \text{if } d \in \{l + 1, \dots, u\}
 \end{array} \right.
 \end{aligned} \tag{2.2}$$

2.4 Analytical Results

This section demonstrates that the optimal decision policy for the ordinal MDP is an index policy. The existence of an index policy is useful for multiple reasons. First, the index policy reduces the computational complexity of the problem, improving tractability for more complex scheduling paradigms. This is particularly meaningful for chronic conditions that can have a wide range of treatment intervals among patients. Second, the index policy offers decision makers insight beyond the optimal action to take in this period. Specifically, calculating an index value for each action allows clinicians to compare multiple actions that may have similar values. For example, if a “short” treatment interval and a “long” treatment interval have similar index values, a more risk-tolerant clinician might be willing to choose the longer treatment interval in order to gain knowledge about

the patient, even if it is not the optimal action. In contrast, a more risk-adverse clinician might decide that gaining more knowledge about the individual patient is not warranted and choose the shorter action. An index policy essentially offers decision makers a menu of treatment options. Utilizing this menu can enable joint decisions between both patients and clinicians when determining treatment regimens. Third, the index policy allows us to determine under which conditions commonly used clinical scheduling policies (e.g. a treat-and-extend policy) are optimal policies. Fourth, the index policy allows us to generate interpretable and generalizable policies for clinical scenarios in which the clinician does not have any prior knowledge about a patient’s health state. We discuss these benefits in greater detail in Section 2.5.

2.4.1 Calculating the optimal policy

We prove that the optimal policy is an index policy by showing that the cost-to-go value $V(P, l, u)$ can be decomposed into several nested sub-problems, each of which can be calculated offline. First, Lemma 2.4.1 shows how to calculate the cost-to-go for any state space with a single available action (e.g. $V(P, l, l)$). Second, Lemma 2.4.2 extends this idea for any state space with n available actions. Third, we show in Theorem 2.4.1 that the cost-to-go for taking any action can be calculated as a function of the single-state cost-to-go values (i.e. the $V(P, i, i)$ values, for all $a_i \in A$). Figure 2.3 provides a visual representation of this idea.

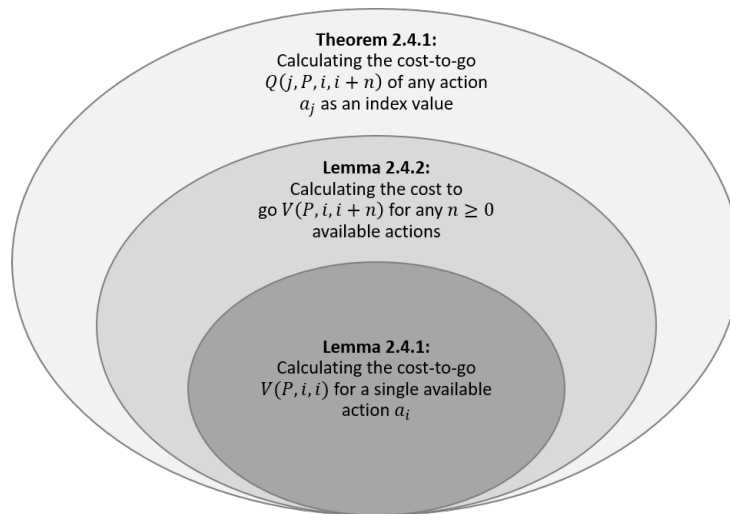


Figure 2.3: A visual representation of the proof for Theorem 2.4.1. We use Lemma 2.4.1 to show how to calculate the cost-to-go for state spaces with a single available action. We then show that the cost-to-go of any action can be calculated as a function of these single-state cost-to-go values $V(P, i, i)$.

Lemma 2.4.1 *The cost-to-go for a single action state space $V(P, l, l)$ can be calculated offline as $V(P, l, l) = \frac{r(l|l)}{1-\delta^l}$.*

A state space with only a single potential action has a single potential decision policy, which is to select that action repeatedly. From Equation 2.2, we have that

$$\begin{aligned}
V(P, l, l) &= \max \left\{ \sum_{i=l}^l (r(l|i) \cdot \phi(i|P, l, l)) + \delta^l V(P, l, l) \right\} \\
&= r(l|l) \cdot \phi(l|P, l, l) + \delta^l V(P, l, l) \\
&= r(l|l) \cdot 1 + \delta^l V(P, l, l) \\
\implies V(P, l, l) &= \frac{r(l|l)}{1 - \delta^l}
\end{aligned} \tag{2.3}$$

■

Lemma 2.4.2 *For any $n \in \mathbb{N}$, the cost-to-go for an $(n + 1)$ -action state space $V(P, l, l + n)$ can be calculated offline.*

We prove this using induction on n . The case of a single action (i.e. $n = 0$) is immediate from Lemma 2.4.1. We next consider the case where there are two possible actions (i.e. $n = 1$). Here, the value function has the form:

$$V(P, l, l + 1) = \max \left\{ \begin{aligned} &\sum_{i=l}^{l+1} (r(l|i) \cdot \phi(i|P, l, l + 1)) + \delta^l V(P, l, l + 1), \\ &\sum_{i=l}^{l+1} (r(l + 1|i) \cdot \phi(i|P, l, l + 1)) + \dots \\ &\delta^{a_{l+1}} \left(\sum_{i=l}^l (\phi(i|P, l, l)) \cdot V(P, l, l) + \sum_{i=l+1}^{l+1} (\phi(i|P, l, l + 1)) \cdot V(P, l + 1, l + 1) \right) \end{aligned} \right\} \tag{2.4}$$

As this is an infinite horizon Markov chain, we know that the optimal policy is a stationary policy (see [75] for more details). The first component of Equation 2.4 shows us that if the decision maker chooses the shortest action a_l , the state space will necessarily remain $(P, l, l + 1)$ in the next decision period. This implies that choosing the shortest action is only optimal in this period if it is also optimal to choose the shortest action forever, since the state space would not change. Then, the resulting cost-to-go for this action is a geometric sum, similar to the single action case. We can substitute this geometric sum into our value function:

$$V(P, l, l+1) = \max \left\{ \begin{array}{l} \frac{\sum_{i=l}^{l+1} (r(l|i) \cdot \phi(i|P, l, l+1))}{1 - \delta^{a_l}}, \\ \sum_{i=l}^{l+1} (r(l+1|i) \cdot \phi(i|P, l, l+1)) + \dots \\ \delta^{a_{l+1}} \left(\sum_{i=l}^l (\phi(i|P, l, l)) \cdot V(P, l, l) + \sum_{i=l+1}^{l+1} (\phi(i|P, l, l+1)) \cdot V(P, l+1, l+1) \right) \end{array} \right. \quad (2.5)$$

We next note that when choosing the longer action a_{l+1} , the state space will necessarily transition into one of two possible single-action states (P, l, l) or $(P, l+1, l+1)$. From Lemma 2.4.1, we know that there is a closed form solution for any single-action state. Therefore, we can substitute these closed form solutions into our two-action equation, and we are then able to calculate the cost-to-go for both possible actions.

We next show the inductive step. Assume that we can calculate $V(P, l, l), V(P, l, l+1), \dots, V(P, l, l+n-1)$ offline. For the $(n+1)$ -action case, we have:

$$V(P, l, l+n) = \max \left\{ \begin{array}{l} \sum_{i=l}^{l+n} (r(l|i) \cdot \phi(i|P, l, l+n)) + \delta^{a_l} V(P, l, l+n), \\ \sum_{i=l}^{l+n} (r(l+1|i) \cdot \phi(i|P, l, l+n)) + \dots \\ \delta^{a_{l+1}} \left(\sum_{i=l}^l (\phi(i|P, l, l)) \cdot V(P, l, l) + \sum_{i=l+1}^{l+n} (\phi(i|P, l+1, l+n)) \cdot V(P, l+1, l+n) \right), \\ \vdots \\ \sum_{i=l}^{l+n} (r(l+n|i) \cdot \phi(i|P, l, l+n)) + \dots \\ \delta^{a_{l+n}} \left(\sum_{i=l}^{l+n-1} (\phi(i|P, l, l+n-1)) \cdot V(P, l, l+n-1) + \dots \right. \\ \left. \sum_{i=l+n}^{l+n} (\phi(i|P, l+n, l+n)) \cdot V(P, l+n, l+n) \right) \end{array} \right. \quad (2.6)$$

We again see that choosing the shortest action a_l transitions to the same state space with a probability of 1. Therefore, the stationary policy principle used in the two-action case also applies here, and the optimal policy will only choose the shortest action in this period if it is always optimal to choose the shortest action in this state. Then, Equation 2.6 can be rewritten as:

$$V(P, l, l+n) = \max \left\{ \begin{array}{l} \frac{\sum_{i=l}^{l+n} (r(l|i) \cdot \phi(i|P, l, l+n))}{1 - \delta^{a_l}}, \\ \sum_{i=l}^{l+n} (r(l+1|i) \cdot \phi(i|P, l, l+n)) + \dots \\ \delta^{a_{l+1}} \left(\sum_{i=l}^l (\phi(i|P, l, l)) \cdot V(P, l, l) + \sum_{i=l+1}^{l+n} (\phi(i|P, l+1, l+n)) \cdot V(P, l+1, l+n) \right), \\ \vdots \\ \sum_{i=l}^{l+n} (r(l+n|i) \cdot \phi(i|P, l, l+n)) + \dots \\ \delta^{a_{l+n}} \left(\sum_{i=l}^{l+n-1} (\phi(i|P, l, l+n-1)) \cdot V(P, l, l+n-1) + \dots \right. \\ \left. \sum_{i=l+n}^{l+n} (\phi(i|P, l+n, l+n)) \cdot V(P, l+n, l+n) \right) \end{array} \right\}, \quad (2.7)$$

We also note that for any action other than the shortest action, the problem will transition into a state with fewer available actions. From our inductive hypothesis, we know that there is a closed-form solution for any of these cost-to-go functions. Thus we can substitute in these closed-form solutions to calculate the cost-to-go of an $(n+1)$ -action space. ■

Theorem 2.4.1 *The optimal decision policy is an index policy, with the index value for action $a_d = Q(d, P, l, u)$ and $V(P, l, u) = \max_{d \in \{l, \dots, u\}} Q(d, P, l, u)$.*

Lemmas 2.4.1 and 2.4.2 demonstrate how to calculate the cost-to-go for any state space with $n \in \mathbb{N}^+$ available actions. In particular, Lemma 2.4.2 shows us that $V(P, l, l+n)$ is a function of the cost-to-go-values of the “nested” state spaces $V(P, i, j)$ where $l \leq i \leq j \leq l+n$. From Equation 2.7, we can define a function $Q(d, P, l, l+n)$ to represent the cost-to-go associated with choosing action a_d in state $(P, l, l+n)$.

$$Q(d, P, l, l+n) := \begin{cases} \frac{\sum_{i=l}^{l+n} (r(l|i) * \phi(i|P, l, u))}{1 - \delta^a} & \text{if } a_d = l \\ \sum_{i=l}^{l+n} (r(d|i) * \phi(i|P, l, l+n)) + \dots & \\ \delta^{a_d} \left(\sum_{i=l}^{d-1} (\phi(i|P, l, l+n)) * V(P, l, d-1) + \dots & \text{if } a_d \in \{l+1, \dots, l+n\} \right. \\ \left. \sum_{i=a}^{l+n} (\phi(i|P, l, l+n)) * V(P, d, l+n) \right) & \end{cases} \quad (2.8)$$

Setting $u = l + n$, we have that $Q(d, P, l, u)$ is the cost-to-go of selecting action a_d . Then, $Q(d, P, l, u)$ is the index value for any action $a_d \in A$ which can be calculated offline, and

$$V(P, l, u) := \max_{d \in \{l, \dots, u\}} Q(d, P, l, u). \quad (2.9)$$

Algorithm 1 in Appendix A provides a method for calculating the index values. ■

2.4.2 Identifying the MSTI

Because $V(P, l, u)$ is based on a set of beliefs P , it is useful to dichotomize the cost-to-go as either 1) the “exploration” component (i.e. decision making to optimize information gained regarding a_M) or, 2) the “exploitation” component (i.e. decision making to optimize key patient health metrics, healthcare costs, etc.). For example, while selecting a treatment interval longer than the MSTI may have short-term unwanted effects to the patient, it also provides valuable knowledge on the true length of the MSTI. This dichotomy helps us understand when the belief-based decision policy will discover the patient’s true MSTI.

As the ordinal MDP incorporates both exploration and exploitation reward types, it can help decisions makers take or avoid calculated risks based on their own values and beliefs. However, as ours is a belief-based model, it is possible that the exploitation component of the rewards may significantly outweigh the exploration component. As a result, the model may not have enough incentive to continue the search for the true MSTI. In Theorem 2.4.2, we provide a condition under which the optimal policy will always find the patient’s true MSTI. The ability to derive such a condition is useful in clinical settings where finding the true MSTI is considered necessary, or at least very important. Theorem 2.4.2 and its Corollary 2.4.1 provide insights into the appropriate model parameters (i.e. the reward function $r(i|j)$ and the discount factor δ) to ensure this condition if applicable.

Theorem 2.4.2 Consider the following reward structure:

- If a patient is not scheduled for treatment and does not have symptoms, they receive a reward $r \geq 0$ in that period.
- If a patient is not scheduled for treatment and does have symptoms, they pay a penalty $c > 0$ in that period.
- If a patient is scheduled for treatment, they receive no reward or penalty (or, they receive a reward of 0) in that period.

Then, there exists a per-period reward to cost ratio r/c above which the ordinal MDP will always converge to the patient's MSTI.

The key idea behind this proof is centered around the condition where the ordinal MDP chooses not to explore longer intervals. From Lemma 2.4.2, recall that if the decision maker selects the shortest available action a_l , the state space will remain the same in the next decision epoch. Thus, if it is optimal to select the shortest action in the current period, it will be optimal to select the shortest action in all future periods due to the stationary policy principle. It follows that if both 1) the optimal policy is to select the shortest available action a_l and 2) action $a_l < a_M$ (i.e. a_l is shorter than the MSTI), then the MDP will never find the MSTI. In other words, the decision maker always gains more information about the MSTI when the shortest possible action a_l does not have the largest index value. Mathematically this is true if there is some action $a_d \in A$ where

$$Q(d, P, l, u) > Q(l, P, l, u) \implies Q(d, P, l, u) - Q(l, P, l, u) > 0 \quad (2.10)$$

Per Theorem 2.4.1, we are able to calculate $Q(d, P, i, j)$ offline for any d , and thus we can check the truth of Equation 2.10 using Algorithm 1. As long as Equation 2.10 holds for all combinations of (l, u) containing the index of the optimal action a_M , the algorithm will converge to the optimal action.

Because of the complex form of the $Q(d, P, l, u)$ values, it is helpful to derive an interpretable form of the threshold in Equation 2.10. As such, we define a lower bound for $Q(d, P, l, u)$ which we write as $Q^-(d, P, l, u)$, where

$$Q^-(d, P, l, u) = \begin{cases} \frac{\sum_{i=l}^u (r(l|i) \cdot \phi(i|P, l, u))}{1 - \delta^{a_l}} & \text{if } d = l \\ \sum_{i=l}^u (r(d|i) \cdot \phi(i|P, l, u)) + \dots \\ \delta^{a_d} \left(\sum_{i=l}^{d-1} (\phi(i|P, l, u)) \cdot Q(l, P, l, d-1) + \dots \right. \\ \left. \sum_{i=d}^u (\phi(i|P, l, u)) \cdot Q(d, P, d, u) \right) & \text{if } d \in \{l+1, \dots, u\} \end{cases} \quad (2.11)$$

The difference between Q and Q^- is that Q^- replaces all cost-to-go functions $V(P, i, j)$ with the index value of the lowest action $Q(i, P, i, j)$. Because $Q(i, P, i, j) \leq V(P, i, j), \forall i, j$ by definition, we know that $Q^-(d, P, l, u) \leq Q(d, P, l, u)$. Then,

$$Q^-(d, P, l, u) - Q(l, P, l, u) > 0 \implies Q(d, P, l, u) - Q(l, P, l, u) > 0 \quad (2.12)$$

Using the argument in Lemma 2.4.2, we can also calculate $Q^-(d, P, l, u)$ for any $d \in \{l, \dots, u\}$. As long as Equation 2.12 is true for all $a_d \leq a_M, d \in \{l, \dots, u\}$, the MDP will eventually select the true MSTI. We also note that Equation 2.12 is equivalent to Equation 2.10 when there are only two potential actions in the state space (i.e. when $u = l + 1$).

To gain insight into the convergence threshold, we consider a common reward structure for $r(d|M)$. Assume that there is a stationary reward $r \geq 0, r \in \mathbb{R}$ for every time period ‘‘at home’’ (i.e. not visiting the clinic) in which no adverse event has occurred, and a stationary cost $c > 0, c \in \mathbb{R}$ for every time period at home in which an adverse event has occurred and has not been treated. For periods in which the patient visits the clinic to receive treatment, there is no reward or cost. For integer values $a_d, a_M \geq 1$, we can write this as:

$$r(d|M) = \begin{cases} 0 & \text{if } a_d = 1 \\ \sum_{i=1}^{d-1} \delta^{a_i} r & \text{if } 1 < a_d \leq a_M \\ \sum_{i=1}^{M-1} \delta^{a_i} r - \sum_{i=M}^{d-1} \delta^{a_i} c & \text{if } a_d > a_M \end{cases} \quad (2.13)$$

Since only one action must have a higher index value than action a_l for convergence, a sufficient condition is to check if choosing action a_{l+1} is preferable to choosing action a_l . Mathematically, we can write this as $Q^-(l+1, P, l, u) > Q(l, P, l, u)$. Under our per-period reward paradigm, it is

useful to write $Q^-(l+1, P, l, u)$ and $Q(l, P, l, u)$ in terms of the probability of an adverse event. Letting $d = l + 1$, Equation 2.12 can be written as:

$$\phi(l|P, l, u) \left(\sum_{i=1}^{l-1} \delta^{a_i} r - \delta^l c + \delta^{a_{l+1}} \frac{\sum_{i=1}^{l-1} \delta^{a_i} r}{1-\delta^{a_l}} - \frac{\sum_{i=1}^{l-1} \delta^{a_i} r}{1-\delta^{a_l}} \right) + (1 - \phi(l|P, l, u)) \left(\frac{\sum_{i=1}^l \delta^{a_i} r}{1-\delta^{a_{l+1}}} - \frac{\sum_{i=1}^{l-1} \delta^{a_i} r}{1-\delta^{a_l}} \right) > 0 \quad (2.14)$$

We can rewrite this as an inequality regarding the reward to cost ratio r/c :

$$\begin{aligned} Q^-(l+1, P, l, u) - Q(l, P, l, u) &> 0 \\ \iff \frac{r}{c} &> \frac{\phi(l|P, l, u) \delta^{a_l}}{\phi(l|P, l, u) \left(\sum_{i=1}^{l-1} \delta^{a_i} + \delta^{a_{l+1}} \frac{\sum_{i=1}^{l-1} \delta^{a_i}}{1-\delta^{a_l}} - \frac{\sum_{i=1}^{l-1} \delta^{a_i}}{1-\delta^{a_l}} \right) + (1 - \phi(l|P, l, u)) \left(\frac{\sum_{i=1}^l \delta^{a_i}}{1-\delta^{a_{l+1}}} - \frac{\sum_{i=1}^{l-1} \delta^{a_i}}{1-\delta^{a_l}} \right)} \\ &= \frac{\phi(l|P, l, u) \delta^{a_l}}{\phi(l|P, l, u) \left((\delta^{a_{l+1}} - \delta^{a_l}) \sum_{i=1}^{l-1} \delta^{a_i} \right) + (1 - \phi(l|P, l, u)) \left(\frac{\sum_{i=1}^l \delta^{a_i}}{1-\delta^{a_{l+1}}} - \frac{\sum_{i=1}^{l-1} \delta^{a_i}}{1-\delta^{a_l}} \right)} \end{aligned} \quad (2.15)$$

■

The threshold in Equation 2.15 offers insight into the requirements for convergence. The numerator of this threshold corresponds to the likelihood of experiencing an adverse event between now and the next visit given action a_d and the current state (P, l, u) . The denominator has two primary components: 1) the expected loss in long-term rewards if the decision maker takes a risk and the patient experiences an adverse event and, 2) the long-term expected value of potentially learning of a new safe interval that is longer than l . The relationship shown in Equation 2.15 matches clinical intuition. If the current belief state suggests a high likelihood of experiencing an adverse event through longer scheduling intervals, the per period value of rewards r must be higher to take the risks associated with finding the MSTI.

From Equation 2.15, we can also see that the discount factor δ plays an important role in the threshold value. In particular, we have the following corollary:

Corollary 2.4.1 *Under the per-period reward structure with $r, c > 0$, the reward to cost threshold goes to 0 as the discount factor δ approaches 1.*

Because only one action needs to have a higher index value than action a_l , it is again sufficient to show that this is true for the threshold shown in Equation 2.15, where $d = l + 1$. We can take the limit of the threshold in Equation 2.15 as δ approaches 1 from the left:

$$\begin{aligned}
& \lim_{\delta \rightarrow -1} \left(\frac{\phi(l|P, l, u)\delta^{a_l}}{\phi(l|P, l, u) \left((\delta^{a_{l+1}} - \delta^{a_l}) \sum_{i=1}^{a_{l-1}} \delta^{a_i} \right) + (1 - \phi(l|P, l, u)) \left(\frac{\sum_{i=1}^l \delta^{a_i}}{1 - \delta^{a_{l+1}}} - \frac{\sum_{i=1}^{l-1} \delta^{a_i}}{1 - \delta^{a_l}} \right)} \right) \\
&= \frac{\phi(l|P, l, u)}{\phi(l|P, l, u) \left(\frac{1}{a_l} - 1 \right) + (1 - \phi(l|P, l, u))(\infty)} = 0
\end{aligned}$$

■

Corollary 2.4.1 suggests that as the importance of the future increases, it is always worth selecting a longer decision to gain more information about the patient’s MSTI. This is an important consideration for clinical contexts in which finding the true MSTI is important for health or financial reasons. It suggests that a sufficiently large discount factor can help ensure that patients eventually reach their true MSTIs while still taking advantage of the ordinal MDP’s suggested exploration policy.

2.4.3 Relaxing the stationary MSTI assumption

For conditions like AMD, each patient’s MSTI is generally stationary throughout a patient’s treatment. [57, 39, 22] However, we can still generalize our model to situations in which the MSTI may change over time to gain insights for other chronic conditions. There are several approaches to this generalization. One example is instead of assuming that a given treatment interval will or will not cause fluid, we can assign a probability distribution to each interval representing the likelihood of fluid when choosing that interval. Then, we can update these distributions over time based on the patient’s measured response to each selected treatment interval. In practice, this distributional approach is not very valuable. The distributional convergence is typically much slower than even the treat-and-extend approach, and can often suggest that clinicians test intervals that cause symptoms multiple times.

Another approach with better practical outcomes is to assume that the MSTI is stationary throughout the initial exploration component. Then, once a patient has reached a steady state, there is some time-to-event probability representing the time until the MSTI changes. This allows for the exploration speed necessary for good clinical outcomes while still generalizing the model for potential MSTI changes. This is the approach that we focus on in this section.

Under our modeling approach, once a patient reaches a steady state scheduling paradigm (i.e. is being treated at their MSTI) it will be obvious if the MSTI ever becomes shorter. For example,

if the patient is being regularly treated at 6 weeks intervals, if the MSTI reduces to 5 weeks, the patient will show up at the clinic with symptoms, and this will immediately flag that the MSTI has gotten shorter. The more challenging situation is determining when the MSTI has increased. The clinician will only be able to observe an MSTI increase if they test a longer treatment interval. Testing a longer interval has inherent risk, as it exposes the patient to a treatment interval that was previously shown to cause symptoms. The key question becomes: when is it worth testing a longer interval to see if the MSTI has increased? We answer this question directly in Theorem 2.4.3.

Theorem 2.4.3 *Assume that the time until a patient's MSTI increases occurs via an exponential distribution with rate λ . Then, the clinician should only check for a longer MSTI if it has been at least τ time periods since the last check, where*

$$\tau = \frac{1}{\lambda} \ln \left(\frac{V(P, M + 1, M + 1) + r(M + 1|M) + \delta^{M+1}V(P, M, M)}{V(P, M + 1, M + 1)} \right). \quad (2.16)$$

The most general answer to the question of when to test for an MSTI increase can be answered by the following inequality:

$$P(\text{longer MSTI}) * \text{Benefit}(\text{longer MSTI}) > P(\text{same MSTI}) * \text{Risk}(\text{same MSTI})$$

In words, it is worth testing for a longer MSTI when the expected benefit of treating less often is greater than the expected risk of exposing the patient to symptoms.

For our analysis, assume that a patient's MSTI increases according to a time-based exponential distribution with rate $\lambda \geq 0$. Specifically, let $L(\tau)$ represent the probability that a patient's MSTI increases within τ time periods of finding the original MSTI. Then we have

$$L(\tau) = 1 - e^{-\lambda\tau}. \quad (2.17)$$

We can also quantify the value of testing a longer MSTI given that the MSTI is in fact longer. Assuming that a clinician extends the treatment interval by one period, the long-term value of a longer MSTI is

$$\text{Benefit}(\text{longer MSTI}) = V(P, M + 1, M + 1).$$

In contrast, if the MSTI is in fact the same, the patient will be exposed to one period of symptoms and then revert to being treated at the original MSTI. We can write this as

$$\text{Risk}(\text{same MSTI}) = r(M + 1|M) + \delta^{M+1}V(P, M, M).$$

We can then formalize our testing inequality as the following:

$$L(\tau)V(P, M + 1, M + 1) > (1 - L(\tau))(r(M + 1|M) + \delta^{M+1}V(P, M, M)) \quad (2.18)$$

In order to isolate the τ we can rewrite Equation 2.18 as

$$\begin{aligned} L(\tau)V(P, M + 1, M + 1) &> (1 - L(\tau))(r(M + 1|M) + \delta^{M+1}V(P, M, M)) \\ \implies (1 - e^{-\lambda\tau})V(P, M + 1, M + 1) &> e^{-\lambda\tau}(r(M + 1|M) + \delta^{M+1}V(P, M, M)) \\ \implies e^{-\lambda\tau} &< \frac{V(P, M + 1, M + 1)}{V(P, M + 1, M + 1) + r(M + 1|M) + \delta^{M+1}V(P, M, M)} \\ \implies -\lambda\tau &< \ln\left(\frac{V(P, M + 1, M + 1)}{V(P, M + 1, M + 1) + r(M + 1|M) + \delta^{M+1}V(P, M, M)}\right) \\ \implies \tau &> -\frac{1}{\lambda}\ln\left(\frac{V(P, M + 1, M + 1)}{V(P, M + 1, M + 1) + r(M + 1|M) + \delta^{M+1}V(P, M, M)}\right) \\ \implies \tau &> \frac{1}{\lambda}\ln\left(\frac{V(P, M + 1, M + 1) + r(M + 1|M) + \delta^{M+1}V(P, M, M)}{V(P, M + 1, M + 1)}\right) \end{aligned}$$

Since $r(M + 1|M) + \delta^{M+1}V(P, M, M) \geq 0$, we know that the value of the log is always non-negative. ■

Theorem 2.4.3 provides useful intuition around when to test for a longer MSTI. Note that for an exponential distribution, $\frac{1}{\lambda}$ is the mean time to an event. In this case, $\frac{1}{\lambda}$ represents the mean time until a patient's MSTI increases. Then, Theorem 2.4.3 says that the recommended time to test for a longer MSTI is the average time until an MSTI increase scaled by the log of the relative risk-benefit ratio for the patient. As the cost of symptoms for that patient increases, the time to test increases. In contrast, as the potential benefit of an MSTI increases, the time to test decreases.

2.4.4 Incorporating a patient health state

Conditions managed with a treat-and-extend protocol have a unique feature: clinicians intentionally expose patients to symptoms for a brief period of time in order to determine their MSTI. For these conditions, while these brief exposures to symptoms might have short-term consequences (e.g. blurred central vision in the case of AMD) they typically do not have significant long term effects (e.g. permanent vision loss with AMD). [32, 23, 10] Otherwise, clinicians would be unwilling to expose patients even temporarily to these symptom windows. Even if there is some effect on long-term health, it is generally difficult to quantify during these brief exposures. For this reason,

our baseline model described in Section 2.3 does not incorporate a patient health variable into the state space. Instead, we focus on short-term rewards and costs regarding symptoms and visit costs.

However, it is possible to generalize our baseline model in Section 2.3 to incorporate a long-term patient health state. For conditions where even short exposures could have long-term consequences, this incorporation is vital in decision making. In this section, we describe how to adapt our model to this paradigm while still enabling the index policy described in Theorem 2.4.1.

We can redefine some of the features of the baseline model, as well as define additional model parameters:

- H : the set of available health states, $H = \{h_1, h_2, \dots, h_K\}$, where $K < \infty$.
- $X = (P, l, u, h)$, the current state with the patient's health h added.
- $\phi(i|P, l, u)$: the current belief that action a_i is the MSTI, $P(a_i = a_M|P, l, u)$. Note that we assume that the MSTI is independent of the patient's health.
- $T(i, X)$: the state transition matrix, which now includes the patient health.
- $T(h'|i, X)$: the probability that patient's health state in the following period will be h' given action a_i and current state X .
- $r(i|M, h)$: the reward for selecting action a_i given the MSTI a_M and the patient's health state h .

To align with our assumption that a patient's MSTI is stationary over time, we also assume that a patient's MSTI is independent of their health state. Additionally, we assume that a patient's health state only changes (deteriorates) if they are exposed to symptoms. To account for the dynamically changing health state, our simplified optimality equation in Equation 2.2 must be rewritten to incorporate health transitions. With a health state, our optimality equation becomes:

$$\begin{aligned}
V(X) = & \\
\max_{d \in \{l, \dots, u\}} & \left\{ \begin{array}{ll}
\sum_{i=l}^u (r(l|i, h) \cdot \phi(i|P, l, u)) + \delta^{a_l} V(P, l, u, h), & \text{if } d = l \\
\sum_{i=l}^u (r(d|i, h) \cdot \phi(i|P, l, u)) + \dots \\
\delta^{a_d} \left(\sum_{i=l}^{d-1} (\phi(i|P, l, u)) \cdot V(P, l, d-1, h) + \dots \right. \\
\left. \sum_{i=d}^u \left(\phi(i|P, l, u) \cdot \sum_{h'} (T(h'|i, X) \cdot V(P, d, u, h')) \right) \right), & \text{if } d \in \{l+1, \dots, u\}
\end{array} \right.
\end{aligned} \tag{2.19}$$

With the above incorporation of a patient health state, we use Corollary 2.4.2 to show that using this formulation retains the index policy feature of Theorem 2.4.1.

Corollary 2.4.2 *For the ordinal MDP with a discrete, finite set of health states H , the optimal policy remains an index policy if preventing symptoms also prevents disease progression.*

We prove this using the same approach as Theorem 2.4.1. In Lemma 2.4.1, we show that the value function for a patient with a known MSTI can be calculated offline. That is, we show that $V(P, i, i)$ can be calculated offline for any action $a_i \in A$. To extend this to the current scenario, we must show that we can calculate $V(P, i, i, h)$ offline for all $a_i \in A$ and $h \in H$. Here, because we have narrowed the action space down to a single action, the only component of our state space that changes over time is the patient's health state. For any $a_i \in A$ we can rewrite our optimality equation as

$$V(P, i, i, h) = r(i|i, h) \cdot \phi(i|P, i, i) + \delta^{a_i} \sum_{h'} (T(h'|i, X) V(P, i, i, h')) \quad \forall h \in H \tag{2.20}$$

From Equation 2.20 we have $[H]$ equations with $[H]$ unknowns for each potential action a_i , and can therefore solve for the cost-to-go values directly. Since we are able to calculate the cost-to-go values for all actions with a known state, we can then apply the proof by induction results from Lemma 2.4.2 to see that the index value solution is also true for this health state modeling approach. ■

2.5 Numerical Results

We next evaluate the results derived in Section 2.4 via a case study on neovascular AMD, utilizing real-world patient data collected from the University of Michigan’s Kellogg Eye Center. We first discuss model parameterization using both patient data and clinical expertise. We provide an example of a patient following the ordinal MDP treatment policy, and describe best practices for treating patients whose MSTI is difficult or impossible to estimate prior to starting treatment. We then simulate the use of the ordinal MDP policy on Kellogg patients. We measure key patient metrics and compare them to the outcomes of following the commonly used treat-and-extend policy, and discuss the effects of the reward function and discount factor on model performance. All analyses in this section were performed using Python version 3.7.11. The University of Michigan Institutional Review Board approved the use of patient data in this study (HUM00156855).

2.5.1 Data sources

The data for this study were collected from neovascular AMD patients treated at the Kellogg Eye Center between June 2013 and March 2020. To be included in this study, patients must have received anti-VEGF injections as their primary method of treatment. We excluded patients that stopped receiving treatment before reaching an MSTI. There are many reasons patients may prematurely stop receiving treatment, including switching providers, treatment refusal, or death. We identified an MSTI as 4 or more consecutive treatments with the same amount of time between visits, within 1 week. We excluded patients that did not have archived OCT information at their visits, which are a common diagnostic tool for managing AMD, and are used in our models for generating initial beliefs. [98] A total of 242 patients met these inclusion criteria, and the distribution of MSTIs among patients is shown in Figure 2.4.

2.5.2 Model parameterization

The core parameters for our model are the action space A , the reward function $r(i|j)$, the discount factor δ and the initial beliefs P . All other model parameters can be derived from these four core parameters as shown in Section 2.3. The first three of these are primarily derived from clinical expertise, as discussed below:

- A : We use a set of actions $A = \{4, 5, \dots, 12\}$ weeks. The range and spacing of treatment intervals is based on average patient response to current injection medications as shown in clinical trials and used in the treat-and-extend protocol. [98]

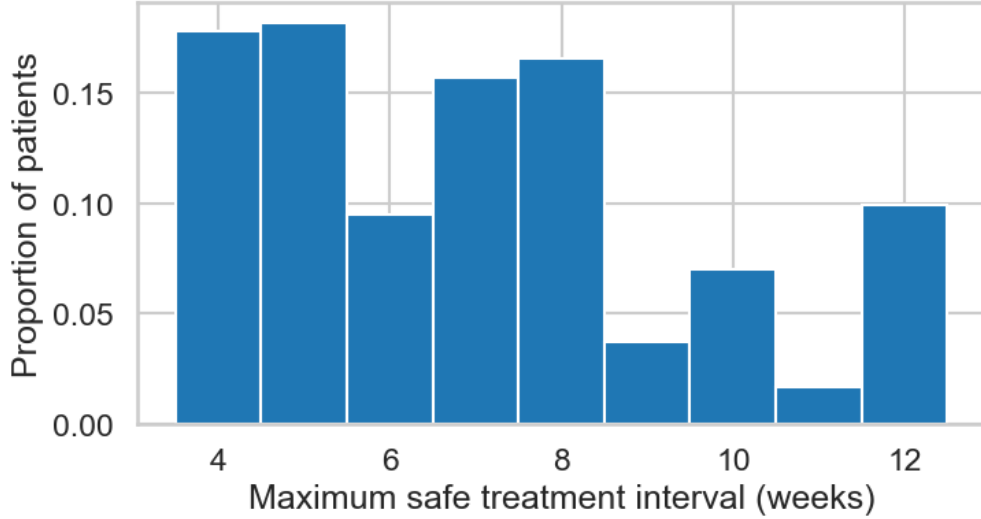


Figure 2.4: The proportion of Kellogg patients with a given maximum safe treatment interval.

- $r(i|j)$: We use the reward function described in Equation 2.13. As a sensitivity analysis, we test a range of reward to cost ratios where $r/c \in [2^{-1}, 2^5]$. To our knowledge, there are no current studies that directly relate a week of time with or without fluid to long-term health metrics, such as quality-adjusted life years. In lieu of this information, we worked closely with our clinical collaborators to discuss the relative importance of a week with no fluid versus a week with fluid on long-term vision when determining the range of our sensitivity analysis. As a result of these discussions, we selected $r/c = 2.0$ as a baseline scenario.
- δ : As a sensitivity analysis, we test a range of discount factors $\delta \in [0.990, 0.999]$. Similar to the reward function, we worked closely with our clinical collaborators to discuss the importance of future reward values. In general, clinicians value the future significantly, as their goal is ultimately find a long-term treatment paradigm for the patient. We selected $\delta = 0.995$ as a baseline scenario.

2.5.2.1 Initial beliefs

For the fourth parameter P , there are a number of methods for deriving the set of initial beliefs. In this paper, we offer two approaches:

- **Using a machine learning predictive model:** One method for predicting the initial beliefs is to take baseline patient health measurements and use a machine learning model to estimate probabilities. In our analysis, we used a multinomial logistic regression, which output the likelihood that each element of A is the true MSTI. The covariates included patient age, sex, visual acuity, and retinal thickness as measured by an OCT scan. To parameterize our


classification model, we fit a model using the `sklearn` package version 0.24.2 in Python on a training subset. Other classification methods such as ordinal regression, neural networks, or support vector classification can also be used to predict P . For our case study, multinomial logistic regression provided the highest predictive power amongst all tested models, with the primary metric being area under the curve (AUC). We used this prediction method as our baseline approach.

- **Using the population level prevalence of MSTIs:** A second method for calculating the initial beliefs is to use the prevalence of MSTIs across the entire patient population. For example, in our patient subset, approximately 18% of patients have a 4-week MSTI. Thus, we can initialize $P(a_M = 4) = 0.18$. The benefit of this approach is that it requires no individual patient health information prior to treatment, and can be used even when such information is unavailable. One potential drawback of this approach is that it uses the same initial beliefs across the population, and ignores potentially insightful patient information that could inform a better policy for individual patients.

Of note, the predictive ability of any machine learning approach to initializing P can vary significantly. For example, the Kellogg Eye Center data available to our team contained only high-level patient information, and the average AUC of our best model was 0.58. As a sensitivity analysis on the effects of this predictive ability we also tested a model built on simulated patient data with an average AUC of 0.70. From related literature, we believe this is a conservative estimate of the potential predictive ability of a machine learning model had we had access to more detailed OCT information (i.e. full images resulting from OCT scans). [11, 86, 85, 78] We generated our simulated data using `sklearn` in Python, and designed it to have a similar size and number of features as the Kellogg patient data. We applied the same multi-class logistic regression model to this data, and added noise to the data in order to reduce the model’s average AUC to 0.70. On the other end of the predictive ability spectrum, the prevalence-based belief analysis ignores all individual patient information and uses only population level information. We therefore use it to represent the scenario with the lowest reasonable predictive ability.

2.5.3 Example patient

To illustrate the ordinal MDP framework, we present an example patient in Figure 2.5. Here, the clinician chooses among weekly scheduling options ranging from 4- to 12-week intervals. The patient has an MSTI of $a_M = 8$ weeks, which is unknown to both the clinician and the model, but can be used to validate the outcome of scheduling decisions. The clinician first observes the patient’s initial health information, including demographics and retinal thickness measurements


MSTI: 8 weeks
Discount factor: 0.995
Reward to cost ratio: 2.0
Actions available: 4,5,..., 12 weeks

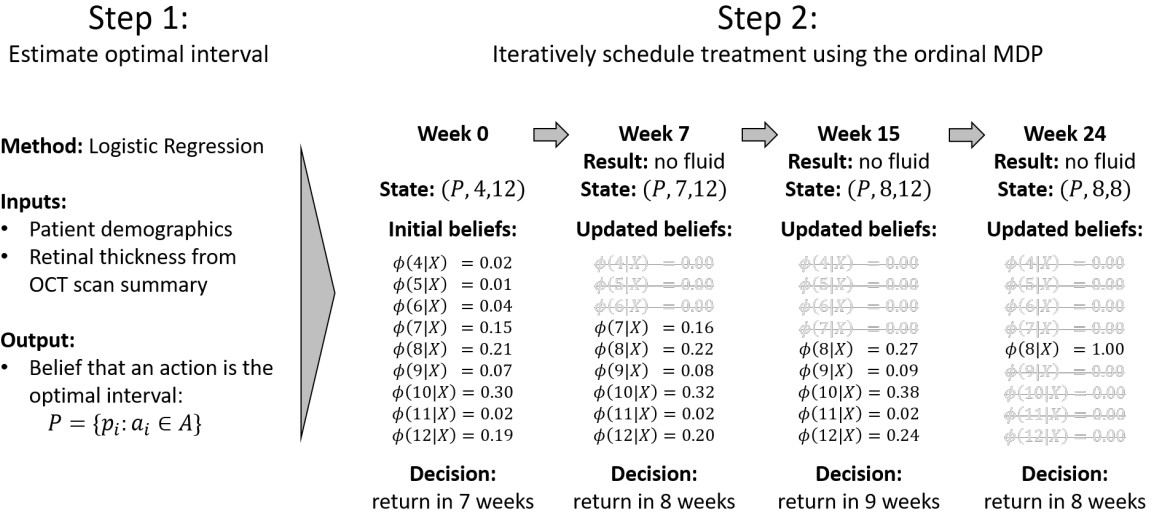


Figure 2.5: An example of a patient following the ordinal MDP using regression-based initial beliefs.

from an OCT scan. This information is input into a multi-class logistic regression model, which estimates the probability set P , representing the likelihood that a given interval is the MSTI. The initial state space X for the patient includes this set P , as well as the lower and upper scheduling intervals for this problem, which are initialized as $a_l = 4$ and $a_u = 12$. We solve for the index values using Algorithm 1 in Appendix A, and the model recommends that the patient return in 7 weeks. Because this decision is less than the true MSTI, there is no fluid when the patient returns. The clinician then updates her state space $X = (P, l, u) = (P, 7, 12)$ and her corresponding belief set $\Phi(X)$. After recalculating the index values, the model recommends the patient return in 8 weeks. As this is the MSTI, there is again no fluid when the patient returns. Still, the model does not yet know with certainty that this is the MSTI, and the state space is updated to be $X = (P, 8, 12)$. Solving for the index values recommends the patient return in 9 weeks, in order to reduce the uncertainty around the MSTI. As 9 weeks is longer than the MSTI, there is fluid when the patient returns, and the state is updated to $X = (P, 8, 8)$. At this point, the only available action is to return every 8 weeks indefinitely, and the decision making process can terminate. In this scenario, the patient experiences 1 week with fluid and requires 3 injections until learning the MSTI.

For context, we can compare the ordinal MDP policy to the treat-and-extend policy. Recall that, after an initial dosing period using 4-week intervals, treat-and-extend explores by scheduling

the patient to return after a 6-week interval. If the patient returns with no symptoms, it continues to extend this interval by 2-week increments until symptoms arise or until the patient reaches the maximum interval of 12 weeks. Then, the policy reverts to the longest tested interval in which the patient had no symptoms. If the patient did not reach the 12-week maximum, the policy then performs a finer search by extending the interval using 1-week increments until symptoms again arise or until the patient would reach an interval known to cause symptoms. At this point, the search is complete, and the clinician indefinitely schedules the patient using the longest interval which does not cause symptoms. Had this patient followed the treat-and-extend policy, the intervals would have been $6 \rightarrow 8 \rightarrow 10 \rightarrow 8 \rightarrow 9$ weeks before treating at 8-week intervals indefinitely. These 5 injections required to learn the optimal interval would have resulted in 3 weeks with fluid.

We can also compare the regression-based initial beliefs to other prediction methods. Had the clinician used the ordinal MDP with prevalence-based initial beliefs, the intervals would have been $5 \rightarrow 6 \rightarrow 7 \rightarrow 8 \rightarrow 9$ weeks before treating at 8-week intervals indefinitely. This would have required 5 injections to learn the MSTI and resulted in 1 week of fluid.

2.5.4 Optimal policies for population-level initial beliefs

For approaches using a common set of initial beliefs, such as the prevalence-based belief method, we can describe general treatment recommendations for all patients. Not varying the initial beliefs P means that all patients follow the same policy based on the current values of l and u . Describing these policies is useful for two reasons: 1) the results are informative in situations in which clinicians cannot make personalized predictions and 2) it allows for direct comparison to (and improvement upon) clinical policies such as treat-and-extend. In Figure 2.6, we show the recommended action to take when following the prevalence-based and treat-and-extend policies under the baseline parameters of $A = \{4, 5, \dots, 12\}$, $r/c = 2.0$, and $\delta = 0.995$. In Figure 2.6, actions highlighted in gray show where the treat-and-extend policy differs from the prevalence-based policy.

Interestingly, the optimal policy when using prevalence-based initial beliefs recommends starting with a 5-week interval and extending the treatment interval by 1-week increments until fluid is observed. This suggests that without individual patient information, the optimal policy should use 1-week extensions instead of the 2-week extensions recommended by treat-and-extend. This indicates that the risk of fluid in the current form of treat-and-extend, as well as the backtracking required to find the MSTI, outweighs any advantages from its quicker initial exploration. This is an important finding from a clinical standpoint—even without personalized initial beliefs, by solving the bandit we can make a data-driven recommendation for improving the current treat-and-extend framework.

Prevalence-based initial beliefs										Treat-and-extend													
		u												u									
		4	5	6	7	8	9	10	11	12			4	5	6	7	8	9	10	11	12		
l	4	4	5	5	5	5	5	5	5	5			4	5	6	6	6	6	6	6	6		
	5		5	6	6	6	6	6	6	6			5	6	6	6	6	6	6	6	6		
	6			6	7	7	7	7	7	7			6		6	7	8	8	8	8	8		
	7				7	8	8	8	8	8			7			7	8	8	8	8	8		
	8					8	9	9	9	9			8				8	9	10	10	10		
	9						9	10	10	10			9					9	10	10	10		
	10							10	11	11			10						10	11	12		
	11								11	12			11							11	12		
	12									12			12								12		

Figure 2.6: The optimal actions to take under the prevalence-based ordinal MDP and the treat-and-extend policies. As the initial beliefs P are shared across all patients, we show the optimal action in terms of the remaining state parameters l and u . Highlighted cells show where the policies differ.

2.5.5 Population level analysis

We simulated the application of the ordinal MDP on our subset of Kellogg Eye Center patients to test the model’s efficacy. Figure 2.7 is a visual representation of our discrete event simulation. After randomly dividing the patient data into a 50/50 training/testing split, the simulation initializes the parameters A , δ , r , and c . If there are I patients in the test set, for each patient i the simulation then estimates P and calculates the index values for every possible combination of l and u values using Algorithm 1. Calculating the index values for all possible l and u combinations provides a complete description of the optimal treatment policy for that patient. Once that policy is established, the simulation applies it to the patient by selecting the highest index value given the current l and u values. Any ties are broken by selecting the shorter of the two intervals. After this selection, the simulation calculates the key patient metrics given the true value of a_M for the current patient.

In Figure 2.7, we represent the previously selected action as a_{-1} . We also use N_v to represent the number of visits until the patient reaches an MSTI and W_f to represent the cumulative number of weeks the patient experiences symptoms (i.e. exudative fluid). While the model does not know the MSTI of any patient when making recommendations, we are able to validate our results by using the MSTI observed in the data for each patient. For example, if we know that the patient has an 8-week MSTI, we can assume that intervals longer than 8 weeks will result in fluid, while intervals shorter than or equal to 8 weeks will not result in fluid. To better understand potential variance of our estimates of population level parameters, we performed the simulation 100 times, each time using a different 50/50 training/testing split on the population. In each iteration, the

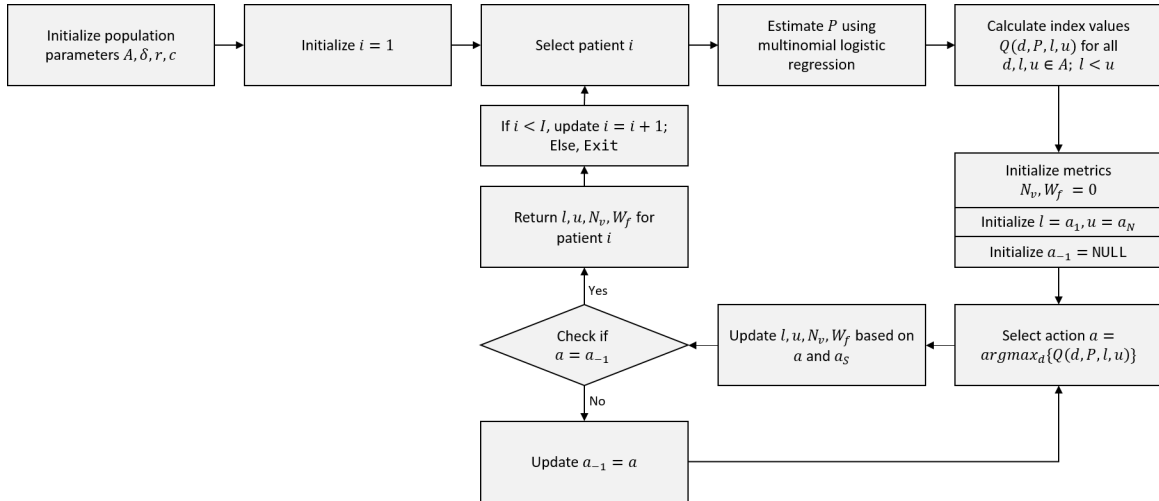


Figure 2.7: A visual representation of the discrete event simulation. To estimate the variance in the optimal interval distribution, the simulation was performed on 100 different 50/50 training/testing splits.

metrics are calculated only using the testing subset of patients.

2.5.6 Population level results

We next discuss the results of the population-level simulation in terms of two key health parameters: 1) the average amount of time patients have fluid in their eye and 2) the average amount of time until the patient reaches a stationary interval. Considering both of these metrics helps capture the balance between finding a patient’s MSTI quickly and not exposing them to unnecessary symptoms. For context, we compare the ordinal MDP results across the different prediction methods described in Section 2.5.2, as well as the treat-and-extend clinical policy.

Figure 2.8 shows the results across a range of reward to cost ratios. For these results, the discount factor is held at the baseline level of $\delta = 0.995$. Figure 2.9 shows the results across a range of discount factors, where the reward to cost ratio is held at the baseline level of $r/c = 2.0$.

Across the entire parameter range studied, following the ordinal MDP policy resulted in fewer weeks with fluid and fewer visits until stationary for patients, regardless of the initial belief method. This suggests that the ordinal MDP is useful even without access to prior patient information or a strong predictive model for the initial beliefs. The prevalence-based model, by using single week extensions starting at 5 weeks, was not only faster than treat-and-extend at finding the MSTI, but also reduced the time with fluid by 38% at baseline. The logistic regression model performed similarly to the prevalence-based model, but was slightly faster by taking larger exploration steps if individual patient information suggested it was appropriate. At baseline, a logistic re-

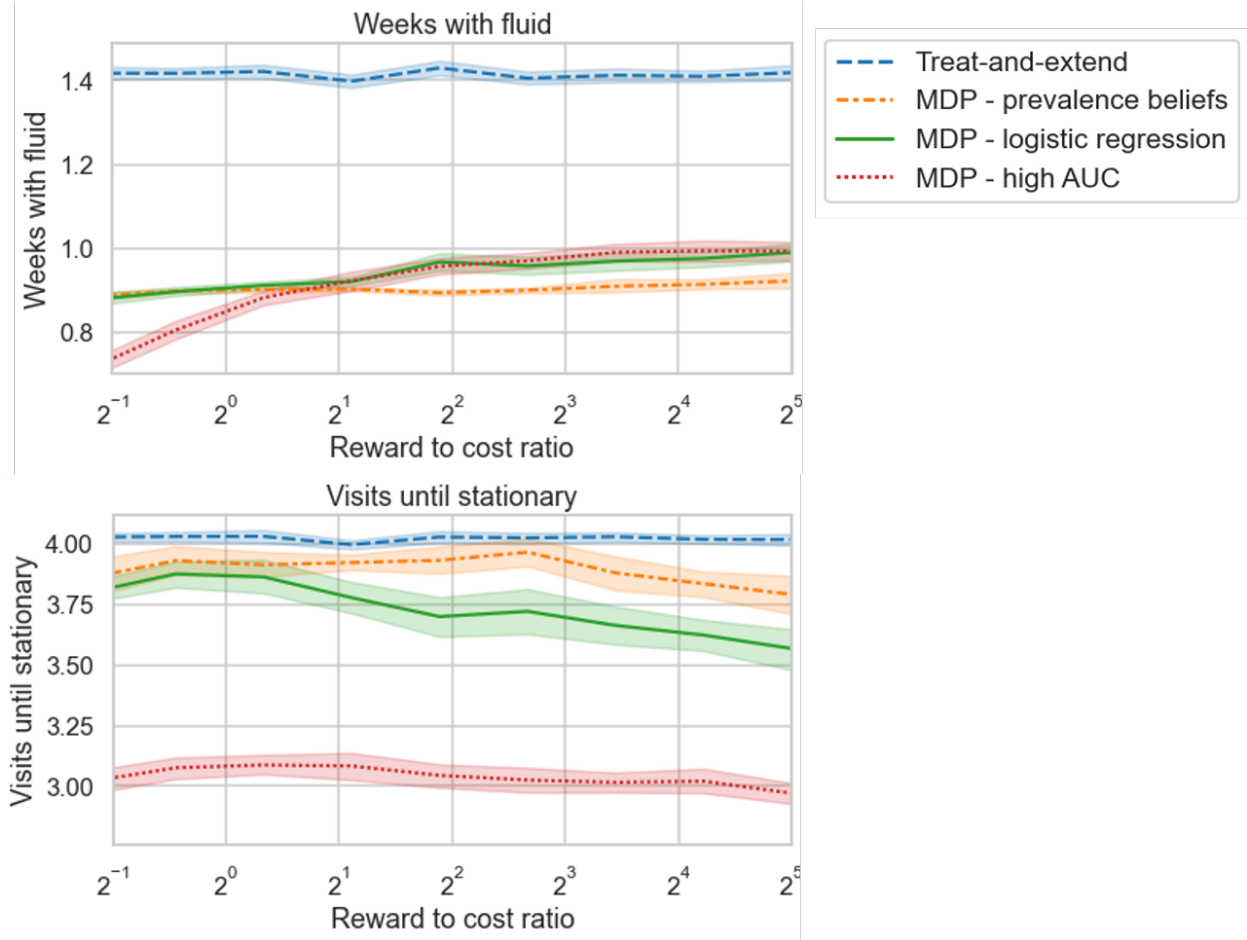


Figure 2.8: Plots comparing the ordinal MDP and treat-and-extend in terms of 1) the average number of weeks with fluid and 2) the average number of visits until reaching a stationary interval. Results are shown across r/c ratios and initial prediction methods. AUC - area under the curve

gression prediction model with a higher AUC (in this case, an average AUC of 0.70) significantly reduced the time to find the stationary interval without increasing exposure to symptoms. The improved predictive ability allowed the model to be more strategic and effective in exploration, making larger extensions when appropriate, and being conservative otherwise.

Generally, a very small reward to cost ratio, such as 2^{-1} resulted in the fewest weeks with fluid. A small r/c suggests that it is much more important to avoid fluid than it is to minimize the number of visits. Consequently, the model often recommends shorter interval extensions during the exploration phase. Recall that from Theorem 2.4.2, it is also possible that the model might recommend scheduling shorter intervals for the entirety of the patient’s follow up, even if it reduces the likelihood that the patient reaches their MSTI. This not only avoids fluid, but also means that patients reach a stationary policy quickly, because the model does not recommend exploring alternatives. A consequence of not reaching the MSTI is that the patient must attend a larger number

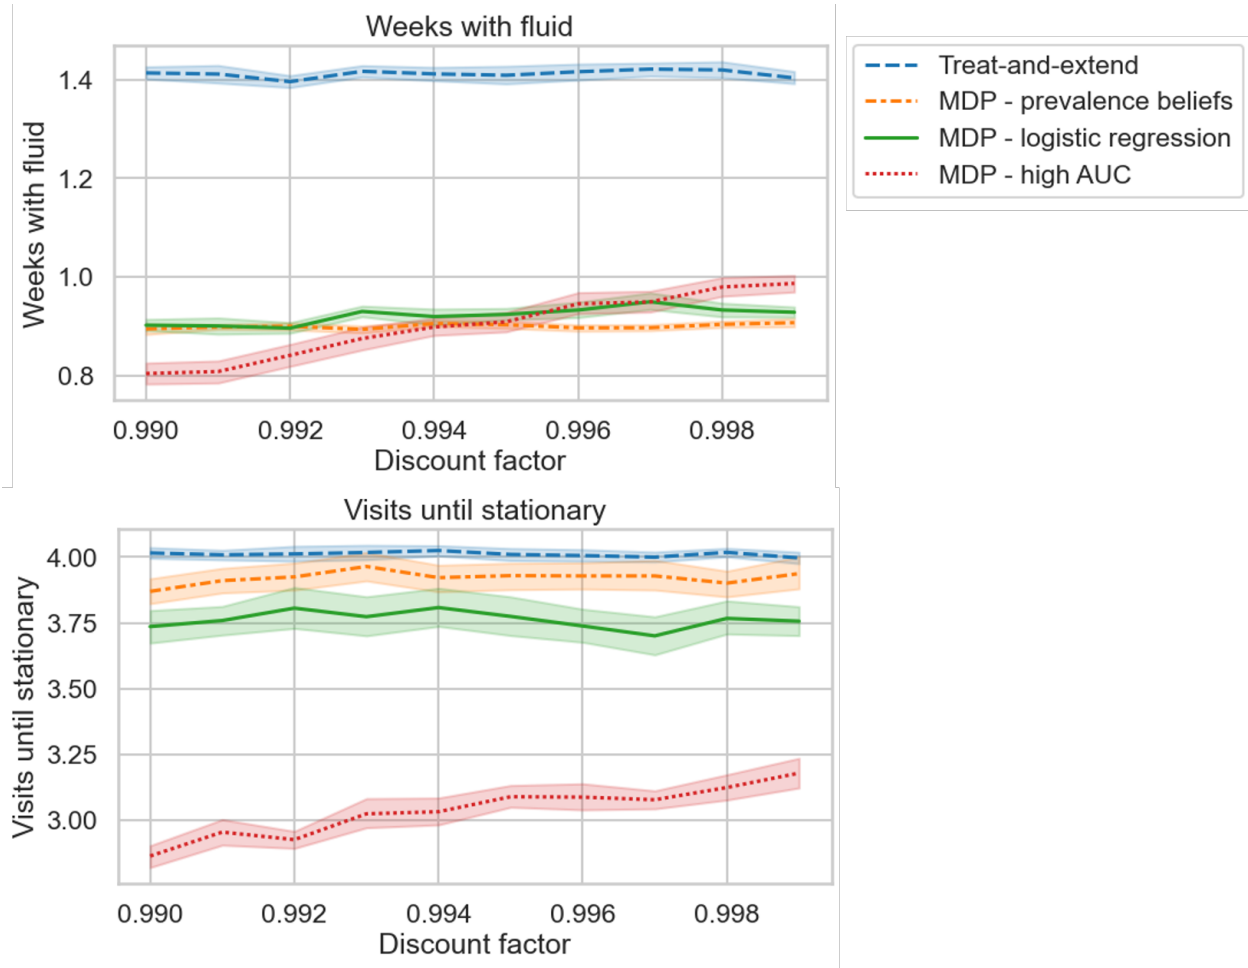


Figure 2.9: Plots comparing the ordinal MDP and treat-and-extend in terms of 1) the average number of weeks with fluid and 2) the average number of visits until reaching a stationary interval. Results are shown across discount factors and initial prediction methods. AUC - area under the curve

of visits in the long run. However, across our baseline reward to cost sensitivity analysis, 100% of patients reached their MSTI, even at the lowest ratio of $r/c = 2^{-1}$. A similar result is seen for discount factors higher than the $\delta = 0.995$ baseline, as suggested by Corollary 2.4.1. The ordinal MDP is designed so that patients will not experience fluid long-term, and any fluid exposure will occur during the exploration phase. As a result, when the future is less important, avoiding fluid becomes more important. However, even at the lowest end of our sensitivity analysis when $\delta = 0.990$, 99.2% of patients reached their MSTI.

In 2018, an estimated 168,400 patients were newly diagnosed with neovascular AMD. [13] At our baseline reward to cost ratio and discount factor, we can extrapolate our average results from Figures 2.8 and 2.9 to a population total. In this scenario, we estimate that following the ordinal MDP would result in approximately 79,100 fewer weeks with fluid and 47,100 fewer injections

overall. Translating a single week with exogenous fluid to a long-term effect on a patient’s vision is difficult, and to the best of our knowledge such a relationship has not been clinically established. However, in 2021, the drug costs for treating AMD ranged from \$50 to \$1,300 per injection, depending on the choice of medication [35], and the average clinical costs were estimated to be \$244 per visit [17]. This suggests that following the ordinal MDP policy when establishing patient treatment schedules could result in a reduction of up to \$61 million in treatment costs annually, in addition to the significant reduction in patient symptoms and long-term costs associated with damaged vision.

2.6 Discussion

In this work we introduce the ordinal MDP as a solution to the fixed-interval treatment planning problem seen in many chronic conditions. In current clinical practice, finding the optimal treatment interval for individual patients is time consuming and can expose patients to unnecessary symptoms and risks. While some clinical heuristics, such as treat-and-extend, currently exist to identify the appropriate treatment interval, they do not leverage the data-driven, optimization-based approach outlined in this paper. The ordinal MDP offers a safe and efficient way to identify a fixed-interval treatment plan for patients diagnosed with a chronic condition.

The ordinal MDP has many practical advantages. First, it can be solved via an index policy, which both reduces the computational complexity typically found in MDP models and provides clinicians with a menu of treatment options. This menu of treatment options is particularly useful in situations in which patients cannot strictly follow an optimal treatment policy due to personal constraints. Second, the ordinal MDP offers insight into the risk-benefit tradeoff found in this exploration problem. Theorem 2.4.2 and Corollary 2.4.1 discuss conditions under which finding a patient’s true steady state interval (i.e. the “exploration” component) is more important than avoiding patient symptoms (i.e. the “exploitation” component). Third, the ordinal MDP allows for the incorporation of personalized health information when designing the optimal policy. By using a machine learning model to establish initial beliefs, decision makers can be strategic in when and how they take risks during the exploration phase. While a higher predictive ability can improve patient outcomes, our case study also suggests that the ordinal MDP is robust to this predictive ability, outperforming current clinical heuristics even without prior patient information.

Beyond its use in scheduling treatment for patients with AMD, we note that the ordinal MDP could be applied to any chronic condition using this fixed-interval treatment paradigm. Outside of treatment scheduling, the ordinal MDP can be generalized to a reward-adjusted binary search. As a result, the model could be applied to several situations where a decision maker’s goal is to guide a process to a steady state, and where the process cannot be observed between decision epochs. As

an example, we can consider a machine maintenance problem where a piece of equipment requires the tuning of a given parameter that can only be adjusted during scheduled downtime. In many cases, key outcomes such as internal wear can only be observed when the equipment is stopped, which comes with significant costs and lost revenue. Balancing the learning curve of this process would be valuable to decision makers and could identify best practices across equipment.

We recognize some key assumptions of the ordinal MDP that limit the scope of this work. First, the baseline formation outlined here requires the assumption that a fixed-interval treatment plan is the appropriate plan, and that an individual patient's optimal interval does not change over time. While we discuss how this assumption can be relaxed in Section 2.4.4, we are unable to parameterize our numerical AMD model using this formulation due to the condition's behavior in clinical practice. Second, the ordinal MDP does not require that the decision maker find a patient's true steady state interval with certainty. As discussed in Theorem 2.4.2 and Corollary 2.4.1, this is a function of the reward function, the discount factor, and the initial beliefs. Should finding the true steady state of all patients be critically important, one easily implementable adaptation of this framework would be to calculate the reward-to-cost threshold that guarantees discovery of a patient's true steady state, and to use a personalized reward-to-cost ratio which is above this threshold. Third, to the best of our knowledge there is no clinically valid method of estimating the effect of a single week of fluid on long-term vision. As a result, we had to restrict our analysis to more indirect metrics, such as time with symptoms. As the clinical understanding of AMD grows, we hope that we can improve our objective function parameters and more directly relate them to primary health metrics.

Overall, the ordinal MDP provides an approach to safely and efficiently identify a personalized fixed-interval treatment schedule for patients. This work builds upon the MAB and MDP literature by introducing an ordinal action space and demonstrates that the optimal policy is an index policy. A simulation on patient data from the Kellogg Eye Center at the University of Michigan illustrates that using the ordinal MDP could reduce patient symptoms while also reducing the number of required visits and overall treatment costs. We hope that our introduction to this framework inspires further development and validation, and eventually serves to improve the treatment of patients diagnosed with chronic conditions.

CHAPTER 3

Synchronizing the Treatment of Multiple Chronic Conditions

3.1 Introduction

For the 4 in 10 U.S. adults diagnosed with multiple chronic conditions, scheduling treatment is a major challenge. [16] Even if the MSTI of each condition is known, the overall optimal treatment schedule is often very unclear. Recall the example shown in Figure 1.2, which highlights a patient with only two chronic conditions. This example patient has MSTIs of 3 and 4 periods, respectively. To minimize treatment costs, the clinician could schedule treatment at the MSTIs. However, by simply scheduling both conditions to be treated every 3 periods, the clinician could synchronize the patient's treatment and avoid several visits to the clinic. A reduced number of visits often means saving the patient significant time and money. However, synchronizing the treatment of the conditions requires treating one of them more frequently than necessary, resulting in additional treatment costs long-term. When coordinating multiple conditions, clinicians must constantly manage this trade-off between minimizing visit costs and minimizing treatment costs.

A common example of managing multiple chronic conditions is again neovascular AMD. Of the 3 million U.S. adults diagnosed with advanced AMD, approximately one-third have bilateral AMD, or AMD in both eyes. [13] The treatment of both of these eyes is managed by a single clinician, meaning that scheduling decisions for each eye are easily synchronized or kept independent. Additionally, an anti-VEGF injection to treat AMD takes only a few minutes, meaning that treating both eyes (versus a single eye) has a marginal impact on the length and cost of a clinical visit. Patients can easily save a trip to the clinic for an injection by synchronizing the injection timing for each eye. However, saving a visit often means additional treatment for one of the patient's eyes, and the cost of treatment can range dramatically depending on which medication the patient is receiving. For example, one commonly used anti-VEGF agent (bevacizumab) costs only \$50 per injection, while the primary alternative medication (aflibercept) costs \$1,600 per injection. [35] The price discrepancy is primarily driven by the fact that bevacizumab was not originally designed

for AMD treatment, but instead for cancer treatment. It was ultimately found to be a viable and often very effective treatment for AMD, and its lower price point is clearly very attractive to patients. [13] However, not all AMD patients respond to bevacizumab. [13] Aflibercept, in contrast, was designed specifically for AMD treatment. However, because of the price structure of these medications, clinicians typically start a patient on bevacizumab treatment to see if they respond well to this medication. If not, clinicians switch the patient to aflibercept despite its higher price point. As might be expected, the decision whether or not to synchronize the injections for both eyes depends heavily on the individual patient's visit and treatment costs.

3.1.1 Chapter outline

In this chapter we answer the following key question: When is it worth synchronizing the treatment of multiple chronic conditions? In Section 3.2, we discuss the current state of the literature on managing multiple conditions. In Section 3.3, we describe our novel dynamic programming approach to treatment scheduling. This model can be used to provide optimal treatment schedules for patients with multiple chronic conditions generally. In Section 3.4 we use our model to demonstrate that there is a visit cost threshold above which the optimal policy is always to synchronize treatment. We also provide the specific synchronization policy in this instance. For the 2-condition case, we provide a closed-form equation representing this threshold. We then describe the effect that treatment costs and the patient's MSTIs have on this threshold. To reduce the computational complexity of the optimal policy for large state spaces, we provide an easily calculable heuristic policy, derived from a small subset of high performing baseline policies. In Section 3.5, we demonstrate the performance of the optimal, baseline, and heuristic policies via a case study on patients with bilateral AMD at the Kellogg eye center. We show that following our framework can reduce direct medical costs by \$1,183 per patient. Extrapolating our findings to a national level, this would translate to \$583 million in savings.

3.2 Relevant Literature

The work in this chapter builds heavily on our work in Chapter 2. As such, much of the literature discussed in Section 2.2 is relevant for this section as well. However, this work is one of the first operations research approaches to managing *multiple* chronic conditions.

From a medical literature perspective, there is a great deal of research illustrating the importance of considering each of a patient's conditions when making medical decisions. [49, 53, 72, 73, 69] Specifically for AMD, several studies have found that synchronized injections is a safe and often cost-effective approach to treatment. [25, 80, 42] Some medical studies have even incorporated the

existence of comorbidities into decision making, such as Taksler et. al (2013), who use a Markov model to estimate how patient behavior changes (e.g. quitting smoking, eating healthier, etc.) effect an individual patient’s life expectancy. [92] However, this and similar studies only consider one-time recommendations for patients, as opposed to the dynamic decision making in this work. [92, 12]

To the best of our knowledge, very few operations research approaches incorporate multiple conditions into their decision making frameworks. Perhaps the closest examples are found in Mason et. al (2014) and Hajjar and Alagoz (2022). [59, 33] Mason et. al (2014) develop an MDP for the simultaneous management of blood pressure and cholesterol. However, their model focuses on when to initiate treatment, and does not provide long-term scheduling recommendations for follow-up treatment. Similarly, Hajjar and Alagoz (2022) enhance their disease screening model by incorporating the health state of a secondary chronic condition. As an example, they show how screening policies for breast cancer can be improved by incorporating whether or not an individual also has type-II diabetes.

Outside of healthcare, there are some parallels between our treatment scheduling problem and a machine maintenance problem. For example, we could re-frame our scenario to represent a technician who must travel to a site to perform preventative maintenance on a set of machines with differing repair needs. The technician may be able to service multiple machines in the same visit, however, there is a cost with each repair. While there are a number of preventative maintenance applications considering multiple machines (e.g. [24, 81, 63]), the reward functions in previous work are insufficient for our scenario. Generally, researchers accept machine breakdowns as a common occurrence. In health applications, a parallel to a “machine breakdown” could have significant long-term effects on patient health, and are avoided whenever possible. This makes applying past machine maintenance work to our problem impractical. Additionally, the current literature does not consider the grouping of repair tasks in the form of site visits over time. In this way, we believe our model could also provide benefit outside of healthcare applications.

In this work, we hope to demonstrate the importance of considering multiple chronic conditions, and set a strong foundation for future research on this topic.

3.3 Modeling Approach

We model this problem using a discrete time, infinite horizon dynamic program. Our decision maker is a clinician simultaneously scheduling the treatment of N chronic conditions. The clinician’s goal is to minimize the lifetime cost of treatment for all conditions without causing patient symptoms. In each time period t the clinician first observes the time since the patient last received treatment for each condition i , represented as $x_{i,t}$. The clinician then decides which, if any, of the




Figure 3.1: The sequence of events in each decision epoch.

conditions to treat in this period. We assume that each condition i has some known maximum safe treatment interval M_i . If $x_{i,t} = M_i$, then the clinician must treat condition i in period t . If the clinician chooses to treat at least one condition, they incur a visit cost $c_0(A)$, where A is an action set representing which conditions to treat. Additionally, administering treatment incurs treatment costs $c(A, X)$, depending on which conditions are being treated and the time since last treatment for each condition. If condition i is treated in period t , then $x_{i,t+1} = 1$, otherwise $x_{i,t+1} = x_{i,t} + 1$. This decision cycle then repeats in all following periods, and is outlined in Figure 3.1. As chronic conditions often affect patients for the remainder of their lifetime, the decision horizon is generally much longer than individual treatment intervals. We therefore model this as an infinite horizon problem, as is common in the medical decision making literature. [4, 87, 91]

3.3.1 Notation

We can describe our model using the following notation:

- t , the current period.
- N , the number of chronic conditions under consideration.
- M_i , the maximum number of periods that condition i can go between treatments, for $i \in [N]$. Without loss of generality, we order the condition indices in terms of increasing M_i values. That is, $i < j \implies M_i \leq M_j$.
- $x_{i,t}$, the number of periods since condition i was treated prior to period t , where $x_{i,t} \in \{1, 2, \dots, M_i\}$. As this is an infinite horizon problem, we eliminate the time subscript when unambiguous.
- X_t , the set of $x_{i,t}$ values for all conditions, $X_t = \{x_{1,t}, \dots, x_{N,t}\}$.
- $a_{i,t}$, the decision whether or not to treat condition i in period t , for $i \in [N]$. If $x_{i,t} < M_i$, then $a_{i,t} \in \{0, 1\}$. If $x_{i,t} \geq M_i$, then $a_{i,t} = 1$ (i.e. the condition must be treated this period).

- A_t , the set of clinician decisions for all conditions in period t , $A_t = \{a_{1,t}, \dots, a_{N,t}\}$.
- $c_0(A)$, the visit cost associated with treating at least one condition. Generally, this represents administrative or overhead costs for the patient or provider (e.g. travel costs, office visit billing costs, etc.). We assume $c_0(\cdot) \geq 0$.
- $c(A, X)$, the treatment cost associated with treatment decisions A and the time since each condition was last treated. Generally, this represents medication costs, as well as long term health effects associated with (not) treating after a given interval. We assume $c(\cdot, \cdot) \geq 0$.
- δ , the per period discount factor, $\delta \in (0, 1)$.

3.3.2 State space and state transitions

We define our state space as the set $X = \{x_1, \dots, x_N\}$, which is the number of periods since condition i was treated prior to the current period, for all $i \in [N]$. The transitions in this model are deterministic, as described in Equation 3.1. If the clinician decides to treat a condition in the current period, the time since last treatment for that condition will be 1 in the following period. If the clinician decides to not treat a condition in the current period, the time since last treatment will increase by 1 in the following period.

$$x_{i,t+1} = \begin{cases} x_{i,t} + 1 & \text{if } a_i = 0 \\ 1 & \text{if } a_i = 1 \end{cases} \quad (3.1)$$

3.3.3 Optimality equation

Our optimality equation is shown in Equation 3.2. We formulate this as a cost-minimization problem. The immediate costs include the visit cost $c_0(A)$ and the treatment costs, $c(A, X)$. We let X' represent the future state in the following period, which is deterministic given X and A .

$$V(X) = \min_A \{c_0(A) + c(A, X) + \delta V(X'|A, X)\} \quad (3.2)$$

3.4 Analytical Results

In this section, we provide analytical insights derived from this multiple condition decision making framework. We discuss scenarios in which synchronizing treatment is beneficial for patients being treated for N conditions. We provide a specific synchronized decision policy that minimizes patient visit costs, and discuss when it is the optimal policy. We also provide a scheduling heuristic

that is useful for complex problem formulations (e.g. when the MSTIs are very large, and thus the state space becomes large) and bound the regret of following this heuristic. For the latter part of this section, we focus our analysis on insights that are useful in the context of treating $N = 2$ conditions, as it is particularly relevant for the treatment of patients with bilateral AMD.

3.4.1 Assumptions

We first outline assumptions which better align our model to the treatment of conditions like neovascular AMD. This allows us to derive more meaningful insights with an immediately practical clinical context. These assumptions are:

- The clinician cannot intentionally cause disease progression. That is, if $x_{i,t} \geq M_i$, then $a_{i,t} = 1$
- There is a fixed visit cost if at least one condition is treated in a decision epoch, and no visit cost otherwise. Mathematically, we say that $c_0(A)$ is equal to a non-negative constant c_0 if any number of the conditions are treated, and is equal to 0 otherwise.
- There is a fixed treatment cost for treating a given condition that does not depend on the time since last treatment. Mathematically, we say that $c_i(1, x_i)$ is equal to a non-negative constant c_i if condition i is treated, and is equal to 0 otherwise.

3.4.2 Synchronizing treatment

An important topic when treating multiple conditions is the potential benefit of coordinating treatment. By scheduling treatment for multiple conditions in the same clinical visit, the clinician can eliminate the need for a additional visits and the associated visit costs. However, coordinating this treatment means potentially over-treating some conditions, which might result in higher treatment costs long-term. In general, it is difficult for clinicians and patients to understand this tradeoff when scheduling treatment. In this section, we describe conditions under which the optimal policy is to always synchronize the treatment of the conditions.

Definition 3.4.1 *A decision policy is a synchronized policy if it never recommends treating a condition $i \in \{2, \dots, N\}$ without also treating condition 1.*

Recall that, without loss of generality, the condition 1 has the shortest maximum safe treatment interval (i.e. $M_1 \leq M_i$). Note that per Definition 3.4.1, condition 1 may be treated on its own as part of a synchronized policy. This is to incorporate situations in which condition 1 requires treatment much more often than condition i . For example, if condition 1 requires weekly treatment

while condition i only requires yearly treatment, it is clearly impractical to treat both conditions at the same frequency.

In this analysis, we will pay special attention to one particular synchronized policy, which we refer to as π_1 . This policy is one which treats condition 1 at its MSTI, and reduces visit costs by treating the other conditions at one of these regular condition 1 visits. We formally define π_1 below, and provide an example patient in Figure 3.2.

Definition 3.4.2 We define a decision policy $\pi_1(X)$ via the following:

- If for any $i \in [N]$ we have $x_i = M_i$ then:
 - $a_i = 1$ and $a_1 = 1$
 - $a_j = 1$ for all $j \neq 1$ s.t. $M_j - x_j < M_1$
 - $a_k = 0$ for all $k \neq 1$ s.t. $M_k - x_k \geq M_1$
- Else:
 - $a_i = 0$ for all $i \in [N]$

In words, π_1 represents the following: if you must treat a condition now (i.e. if $x_i = M_i$), you also treat condition 1 and any other condition that will require treatment before the next condition 1 treatment (i.e. in less than M_1 periods). The π_1 decision policy reduces visit costs by synchronizing treatment to occur only every M_1 periods. In fact, this policy is of particular interest, because it minimizes overall visit costs, which we prove in Lemma 3.4.1 below.

Lemma 3.4.1 For constant visit costs (i.e. when $c_0(A) = c_0 \in \mathbb{R}^+$ for all A), the π_1 decision policy minimizes a patient's lifetime visit costs.

Case 1: Condition 1 will require treatment at least as soon as all other conditions will require treatment (i.e. $M_1 - x_1 \leq M_i - x_i, \forall i \in \{2, \dots, N\}$).

First, we note that scheduling a visit before at least one condition requires treatment is immediately sub-optimal. Such a decision would add extra visit and treatment costs, with no benefit. Therefore, the clinician should not treat until at least one condition requires treatment. In this case, that will include condition 1.

Long term, since the clinician must treat condition 1 every M_1 periods, the minimum number of visits is 1 visit per M_1 periods. Starting from the period in which we first treat condition 1, this means paying c_0 every M_1 periods. The cost of this is an infinite discounted geometric series, which has a lifetime cost of $\frac{c_0}{1-\delta^{M_1}}$. This means that the minimum lifetime visit cost is $\frac{c_0}{1-\delta^{M_1}}$. By

following π_1 , you will only treat on the frequency of M_1 , which means that the long-term visit cost of following π_1 is also $\frac{c_0}{1-\delta^{M_1}}$. Therefore, π_1 achieves the minimum possible visit cost.

Case 2: Condition 1 only requires treatment after at least one other conditions will require treatment (i.e. $\exists i \in \{2, \dots, N\}$ s.t. $M_1 - x_1 > M_i - x_i$).

In this case, if following π_1 , you would wait until one of the conditions (i.e. condition i) requires treatment and then treat at least condition i and condition 1. Beyond this, you would continue to treat every M_1 periods. Alternatively, you could delay treating some of the conditions that do not require treatment at the same time as condition i to potential reduce long-term treatment costs. We next show that this delay (i.e. deviating from π_1) is never optimal.

Should you choose to delay some conditions beyond the condition i treatment visit, you would pay c_0 for each additional delayed visit until condition 1 reaches its MSTI. Once condition 1 reaches its MSTI, we know from Case 1 that the best you can do in terms of visit cost is $\frac{c_0}{1-\delta^{M_1}}$. However, because of the delay, this long-term visit cost is discounted. At best, you only require one additional visit (i.e. one c_0 payment) for the condition i visit before treating condition 1 (assuming no other conditions need to be treated in between, which would only worsen the visit cost). Therefore, the best possible long-term visit cost is $c_0 + \delta^{M_1-x_1} \left(\frac{c_0}{1-\delta^{M_1}} \right)$ when delaying treatment. Note that this cost is calculated starting the period in which you treat condition i . Alternatively, we know that the cost of following π_1 starting in this period is $\frac{c_0}{1-\delta^{M_1}}$. So for deviating to be better than π_1 , we would need:

$$\begin{aligned} c_0 + \delta^{M_1-x_1} \left(\frac{c_0}{1-\delta^{M_1}} \right) &< \frac{c_0}{1-\delta^{M_1}} \\ \implies c_0 &< \left(\frac{c_0}{1-\delta^{M_1}} \right) (1 - \delta^{M_1-x_1}) \\ \implies 1 &< \left(\frac{1 - \delta^{M_1-x_1}}{1 - \delta^{M_1}} \right) \\ \implies \delta^{M_1-x_1} &< \delta^{M_1} \end{aligned}$$

which is a contradiction, since $\delta \in (0, 1)$ and $M_1, x_1 \geq 1$. Therefore, deviating is never optimal, and π_1 also minimizes visit costs in this case. ■

As a consequence of Lemma 3.4.1, we introduce Theorem 3.4.1. This theorem describes a visit cost threshold \bar{c}_0 , above which the optimal policy is a synchronized policy. More specifically, the optimal policy above this threshold is π_1 . This threshold allows us to derive important managerial insights for managing multiple conditions. In the example of AMD, individual patients experience unique visit cost demands. Because AMD can lead to vision loss, a major burden patients experience is coordinating transportation to the clinic for treatment. Patients with significant vision loss may require a caretaker to drive them to the appointment, while others may still be able to

drive themselves. Additionally, rural patients may need to travel long distances to appointments, increasing transportation demands. Long travel times might also require patients to take time off of work to make their appointment, resulting in lost wages. The threshold described in Theorem 3.4.1 helps us understand which categories of patients might require synchronized treatment, and can inform population-level treatment policies.

Theorem 3.4.1 *For constant visit and treatment costs (i.e. if $c_0(A) = c_0 \in \mathbb{R} \forall A$ and $c_i(A, X) = c_i \in \mathbb{R} \forall i, A, X$), there exists a visit cost threshold \bar{c}_0 above which the optimal policy is always π_1 .*

The long-term costs in this decision making problem are a sum of visit and treatment treatment costs. As such, we can write the long-term cost of following any policy as a linear combination of c_0 and c_i values. First, consider the cost of following π_1 , which we can write as:

$$V^{\pi_1}(X) = k_0(X)c_0 + \sum_{i=1}^N k_i(X)c_i$$

In this linear version of our π_1 cost-to-go function, we let $k_0(X)$ represent a visit cost coefficient, and the $k_i(X)$ values represent treatment cost coefficients. Both of these long-term cost scaling coefficients are dependent on the patient's current state X .

Then consider the cost of following some arbitrary decision policy π_2 , which we can similarly write as:

$$V^{\pi_2}(X) = l_0(X)c_0 + \sum_{i=1}^N l_i(X)c_i$$

From Lemma 3.4.1 we have that $k_0(X) \leq l_0(X)$, since π_1 minimizes overall visit costs. We use this fact to show that if c_0 is large enough, π_1 is the optimal decision policy. We do this across two cases.

Case 1: $k_0(X) < l_0(X) \forall X$

In this case, we assume π_1 results in strictly lower overall visit costs. Then, the cost of following

π_1 is at least as good as π_2 if, for all X , we have that

$$\begin{aligned}
& k_0(X)c_0 + \sum_{i=1}^N k_i(X)c_i \leq l_0(X)c_0 + \sum_{i=1}^N l_i(X)c_i \\
\implies & c_0(l_0(X) - k_0(X)) \geq \sum_{i=1}^N (k_i(X) - l_i(X))c_i \\
\implies & c_0 \geq \frac{\sum_{i=1}^N (k_i(X) - l_i(X))c_i}{l_0(X) - k_0(X)}.
\end{aligned}$$

Since $k_0(X) < l_0(X)$, we know that the overall c_0 coefficient is positive and this inequality represents a lower bound on c_0 . Therefore, for any π_2 policy and corresponding $l_0(X)$ and $l_i(X)$ values, there is some finite visit cost value above which π_1 is at least as good as π_2 .

Case 2: $k_0(X) = l_0(X)$ for at least one value of X

In this case, if $k_0(X) = l_0(X)$, then our inequality comparison reduces to

$$\begin{aligned}
& k_0(X)c_0 + \sum_{i=1}^N k_i(X)c_i \leq l_0(X)c_0 + \sum_{i=1}^N l_i(X)c_i \\
\implies & c_0(l_0(X) - k_0(X)) \geq \sum_{i=1}^N (k_i(X) - l_i(X))c_i \\
\implies & 0 \geq \sum_{i=1}^N (k_i(X) - l_i(X))c_i.
\end{aligned}$$

If $k_0(X) = l_0(X)$ then we must have that π_2 also minimizes visit costs. From the proof of Lemma 3.4.1, we know that this can only occur if π_2 also treats every M_1 periods, similar to π_1 . Note that π_1 is also designed to minimize treatment costs while treating on the M_1 frequency (i.e. treat condition i every κM_1 periods where κ is the largest positive integer such that $\kappa M_1 \leq M_i$). Therefore, any alternative policy π_2 which treats on the same frequency as π_1 would have to treat at least one other condition $i \in \{2, \dots, N\}$ more frequently than needed. This means that $k_i(X) < l_i(X)$ for any condition i treated more frequently than π_1 and $k_j(X) = l_j(X)$ for any condition j treated at the same frequency as π_1 . Therefore, the overall treatment costs $\sum_{i=1}^N (k_i(X) - l_i(X))c_i \leq 0$, which means that our inequality comparison always holds, regardless of the value of c_0 . Therefore, π_1 is at least as good as any policy π_2 which fits into this case. ■

In Theorem 3.4.1, we prove the existence of a synchronization threshold \bar{c}_0 . This is useful,

because it tells us that for patients with “high” visit costs, π_1 is always the optimal scheduling policy. From a practical perspective, however, this is less helpful if we cannot calculate the specific value of \bar{c}_0 for an individual patient. In the generic N -condition case, this is a daunting analytical task, because calculating the value of this threshold requires calculating the value of all possible $k_0(X)$, $k_i(X)$, $l_0(X)$, and $l_i(X)$ values. In such cases, it is generally easier to do a numeric search over a range of visit costs to determine when π_1 becomes optimal. However, in Section 3.4.3 below, we provide a closed-form method for calculating these values in the 2-condition scenario. Additionally, we use these formulas to derive key managerial insights for patients being treated for 2 conditions, such as patients being treated for bilateral AMD.

3.4.3 Coordinating two conditions

To get insights regarding the treatment of bilateral AMD, we can rewrite Equation 3.2 in the context of $N = 2$ conditions. In this scenario, there are at most 4 possible actions that the clinician can take in each period: (i) treat both conditions, (ii) treat only condition 1, (iii) treat only condition 2, or (iv) treat neither condition. When all 4 of these actions are available to a clinician, we can reformulate Equation 3.2 as Equation 3.3, with the cost-to-go of all 4 actions listed in the same order as above.

$$V(x_1, x_2) = \min \left\{ \begin{array}{l} c_0 + c_1 + c_2 + \delta V(1, 1), \\ c_0 + c_1 + \delta V(1, x_2 + 1), \\ c_0 + c_2 + \delta V(x_1 + 1, 1), \\ \delta V(x_1 + 1, x_2 + 1) \end{array} \right\} \quad (3.3)$$

Note that these 4 actions are not always available to the clinician. If $x_1 = M_1$, then the clinician must treat at least condition 1 (i.e. actions (iii) and (iv) are not available). Similarly if $x_2 = M_2$, then the clinician must treat at least condition 2 (i.e. actions (ii) and (iv) are not available). In these cases, we can simply write the optimality equation using a subset of the options in Equation 3.3.

In Theorem 3.4.2 are able to derive a closed-form threshold for \bar{c}_0 when $N = 2$ conditions. The threshold described in Theorem 3.4.2 helps us understand specifically which AMD patients would benefit from synchronized treatment, and can inform population-level treatment policies. By deriving this closed form solution, we are also able to gather insights on the effect of problem parameters on the synchronization threshold. For example, the 2-condition threshold can help us understand how medication and administrative pricing affects optimal decision policy. Additionally, Theorem 3.4.2 and it’s corollaries provide insight into how the maximum safe treatment intervals and the discount factor affect any potential advantages of synchronizing treatment.

Theorem 3.4.2 For $N = 2$ conditions, the \bar{c}_0 threshold is:

$$\bar{c}_0 = \max\left\{ \frac{-c_1 * \left(\frac{\delta^{M_1-x_1} - 1}{1 - \delta^{M_1}} \right) - c_2 * \left(\frac{\delta^{(\lfloor \frac{M_2-M_1+x_1}{M_1} \rfloor + 1)M_1-x_1} - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}}{1 - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}} \right)}{\left(\frac{\delta^{M_1-x_1} - \delta^{M_1}}{1 - \delta^{M_1}} \right)}, \right. \quad (3.4)$$

$$\frac{-c_1 * \left(\frac{\delta^{M_1-x_1} - 1}{1 - \delta^{M_1}} \right) - c_2 * \left(\frac{\delta^{M_1-x_1} - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}}{1 - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}} \right)}{\left(\frac{\delta^{M_1-x_1} - \delta^{M_1}}{1 - \delta^{M_1}} \right)},$$

$$-c_1 - c_2 \left(\frac{\delta^{M_2-x_2} - 1}{1 - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}} \right) \left(\frac{1 - \delta^{M_1}}{\delta^{M_2-x_2} - \delta^{M_1}} \right),$$

$$-c_1 - c_2 \left(\frac{\delta^{(\lfloor \frac{M_2-x_2-M_1}{M_1} \rfloor + 1)M_1} - \delta^{(\lfloor \frac{M_2-x_2-M_1+x_1}{M_1} \rfloor + 1)M_1-x_1}}{1 - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}} \right) \left(\frac{1 - \delta^{M_1}}{1 - \delta^{M_1-x_1}} \right),$$

$$-c_1 - c_2 \left(\frac{\delta^{M_1} - \delta^{M_1-x_1}}{1 - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}} \right) \left(\frac{1 - \delta^{M_1}}{1 - \delta^{M_1-x_1}} \right),$$

$$-c_1 - c_2 \left(\frac{\delta^{M_2-x_2} - \delta^{M_1-x_1}}{1 - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}} \right) \left(\frac{1 - \delta^{M_1}}{1 - \delta^{M_1} + \delta^{M_2-x_2} - \delta^{M_1-x_1}} \right),$$

$$\left. \frac{-c_1 * \left(\frac{\delta^{M_1-x_1} - \delta^{M_2-x_2}}{1 - \delta^{M_1}} \right) - c_2 * \left(1 + \frac{\delta^{(\lfloor \frac{M_2-M_1+x_1}{M_1} \rfloor + 1)M_1-x_1} - \delta^{M_2-x_2}}{1 - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}} \right)}{\left(1 + \frac{\delta^{M_1-x_1} - \delta^{M_2-x_2}}{1 - \delta^{M_1}} \right)} \right\}$$

Recall that, π_1 describes a treatment policy where the conditions are synchronized to be treated using intervals that are a multiple of M_1 . Figure 3.2 shows an example of a patient following the π_1 decision policy. In this example, the second condition is treated every 9 periods (as opposed to the MSTI of 11 periods) as a means of reducing visit costs. The consequence of this action is that it increases lifetime treatment costs for condition 2.

To describe the conditions under which π_1 is an optimal policy, we leverage the policy iteration method of solving infinite horizon dynamic programs. Figure 3.3 shows an outline of our approach. We start by calculating the π_1 policy cost-to-go values $V^{\pi_1}(X)$ for all possible values of X . We then perform a single loop of policy iteration. In the policy evaluation step of policy iteration, we




Figure 3.2: An example treatment schedule following π_1 for a patient with $M_1 = 3$ and $M_2 = 11$. Assuming that the patient receives treatment for both conditions at time $t = 0$, this specifically represents the treatment cycle corresponding to the cost-to-go value Φ^{π_1} defined in Appendix B.

first calculate the “would-be” cost-to-go values for each available decision had we selected that action in the current period, and then followed π_1 afterwards. We can define the “would-be” cost-to-go for a given action a as $Q^{\pi_1}(X, a)$. In the policy update step of policy iteration, we update the optimal action $a^*(X) = \arg \max_a \{Q(X, a)\}$.

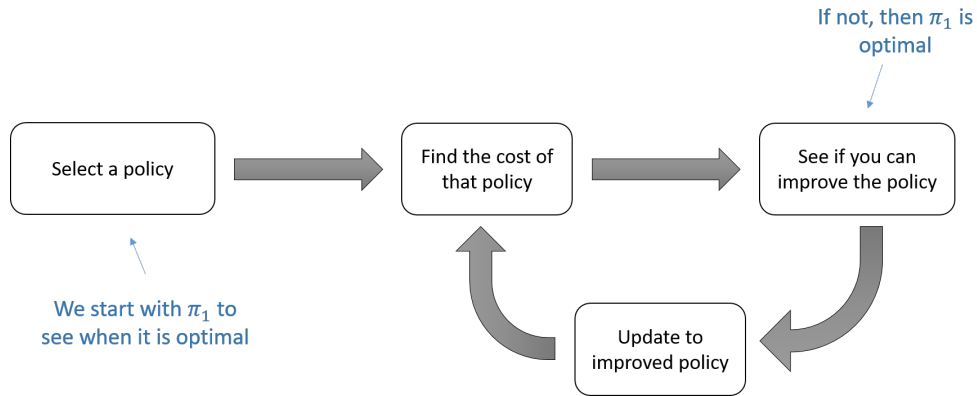


Figure 3.3: An overview of the proof of Theorem 3.4.2. Using the policy iteration method of solving dynamic programs, we select π_1 as our initial policy. We then check whether we can improve upon π_1 , which depends on the problem parameters. The conditions under which we cannot improve upon π_1 are the conditions under which π_1 is the optimal policy.

The key idea of our proof is that π_1 is an optimal policy if and only if $a^*(X)$ matches π_1 exactly for all values of X and the policy iteration algorithm immediately converges. That is, π_1 is optimal if the evaluation step of policy iteration finds that the best (in this case, lowest) cost-to-go values are those that correspond with following π_1 . The number of individual state comparisons makes the proof of Theorem 3.4.2 relatively long, so we put it in Appendix B for brevity. The key result is that there are 7 relevant inequalities across all states in X that must be true for π_1 to be optimal, and these 7 inequalities comprise Theorem 3.4.2.

3.4.3.1 A sufficient synchronization threshold

In Theorem 3.4.2, we derived 7 visit cost thresholds which depend on the current state X of the patient. If the visit cost is above these 7 thresholds for all possible state spaces, then π_1 is universally an optimal policy. We next come to an important observation:

Observation 3.4.1 *In practice, it is unlikely that a patient will visit all possible state spaces.*

Suppose that, at the start of treatment, a clinician treats both eyes at the same period to get patients to a baseline level (i.e. suppose that after the first visit, the patient's state is $X = (1, 1)$). Next, suppose that the patient follows policy π_1 to schedule treatment. As seen in Figure 3.2, the patient will follow a cyclical treatment pattern. As a result, the states that the patient enters will also follow a cyclical pattern, and will consist of a subset of all available states, say S_{π_1} . If the patient begins treatment in some state $X_0 \in S_{\pi_1}$ and follows the π_1 policy, they will never leave this subset of states. Thus, we only need to consider states that the patient will actually enter to determine the conditions under which π_1 is optimal. If π_1 is the optimal policy for every $X \in S_{\pi_1}$, then the patient's state will never leave S_{π_1} , as leaving would be sub-optimal. Therefore, the optimal action in states outside of S_{π_1} do not affect our decision making. If we assume that a patient enters $X = (1, 1)$ after the first visit, we only need to consider the thresholds for states $X \in S_{\pi_1}$.

Definition 3.4.3 *Let S_{π_1} be the set of states visited if following the π_1 decision policy starting from state $X = (1, 1)$. By following the decision rules of π_1 we can describe the set of states in S_{π_1} as*

$$S_{\pi_1} = \{(x_1, x_2) : x_2 = x_1 + lM_1\} \quad (3.5)$$

where l is any integer such that $0 \leq l < \lfloor \frac{M_2}{M_1} \rfloor - 1$.

We can refer to a visit cost threshold that only considers states $X \in S_{\pi_1}$ as a “joint-start threshold”. It represents the threshold required to continue following π_1 if a patient starts in a state within S_{π_1} . Mathematically, let \bar{c}_0 represent the joint-start threshold. As almost all patients begin treatment by getting both eyes to a baseline level, \bar{c}_0 is a practically relevant threshold that will be sufficient for most patients. Since the joint-start threshold \bar{c}_0 includes a fewer number of conditions as the “full threshold” \bar{c}_0 outlined in Theorem 3.4.2, we know that $\bar{c}_0 \leq \bar{c}_0$. In particular, when only considering states within S_{π_1} , we will never enter Case 2 of Theorem 3.4.2. As a result, we can exclude Equations B.4 and B.5 from the conditions required for \bar{c}_0 . The other cases and subcases that can be excluded depend on the patient-specific combination of M_1 and M_2 and are not easily generalizable.

3.4.3.2 The effects of treatment cost on the synchronization threshold

We next consider how the treatment cost parameters c_1 and c_2 affect the synchronization thresholds \bar{c}_0 and $\bar{\bar{c}}_0$. We also show how the joint-start threshold offers important managerial insights that the full threshold does not provide. Specifically, we show the following:

- For joint-start patients, if the cost of treating condition 1 is high enough, π_1 is always optimal, regardless of the visit cost.
- For all patients, if the cost of treating condition 2 is high enough, π_1 is never optimal.
- For all patients, if the treatment cost for both conditions is zero, π_1 is always optimal.
- For all patients, if the MSTI of condition 2 is a multiple of the MSTI of condition 1, π_1 is always optimal.

Corollary 3.4.1 *As the cost to treat condition 1 goes to infinity the full synchronization threshold goes to infinity, but the joint-start threshold goes to zero. That is, as $c_1 \rightarrow \infty$, $\bar{c}_0 \rightarrow \infty$ and $\bar{\bar{c}}_0 \rightarrow 0$.*

To prove this, we look at the coefficients of c_1 across Equations B.4-B.10. In Equations B.4 and B.5, these coefficients are positive. In Equations B.6-B.9, these coefficients are negative. In Equation B.10, this coefficient could be either positive or negative, depending on the state X and the values of M_1 and M_2 . The full threshold \bar{c}_0 includes all of these equations. Since at least two of these equations have positive coefficients, as $c_0 \rightarrow \infty$, the minimum of these equations also goes to infinity. Therefore, \bar{c}_0 also goes to infinity.

Recall that the joint-start threshold $\bar{\bar{c}}_0$ does not include Equations B.4 and B.5, as they are a result of Case 2 of Theorem 3.4.2. We then need to consider the c_1 coefficient in Case 6. When in the subset of states S_{π_1} , if following π_1 , it will always be that $M_1 - x_1 \leq M_2 - x_2$. With this fact, we can see that $\delta^{M_1 - x_1} - \delta^{M_2 - x_2} \geq 0$. Therefore, the coefficient on c_1 will be negative for the states in S_{π_1} . As a result, all of the c_0 coefficients for the relevant joint-start equations are negative. Then, as $c_0 \rightarrow \infty$ the minimum of these equations goes to negative infinity. However, since we assume that $c_0 \geq 0$ we say that $\bar{\bar{c}}_0 \rightarrow 0$. ■

Corollary 3.4.2 *As the cost to treat condition 2 goes to infinity both the full and joint-start synchronization thresholds also go to infinity. That is, as $c_2 \rightarrow \infty$, \bar{c}_0 and $\bar{\bar{c}}_0 \rightarrow \infty$.*

To prove this, it is sufficient to show that one of the equations comprising each of the thresholds goes to infinity as c_2 goes to infinity. Equivalently, we can say that this is true if at least one of the c_2 coefficients is positive. Then, by recognizing that Equation B.8 has a positive c_2 coefficient, and is a relevant threshold for both \bar{c}_0 and $\bar{\bar{c}}_0$, the proof is complete. ■

Corollary 3.4.3 *If the treatment costs for both conditions are 0, the optimal policy is always π_1 . That is, as c_1 and $c_2 \rightarrow 0$, \bar{c}_0 and $\bar{\bar{c}}_0 \rightarrow 0$.*

To show this, we can simply set $c_1 = c_2 = 0$ and see that Equations B.4-B.10 all become $c_0 \geq 0$. ■

Corollaries 3.4.1, 3.4.2, and 3.4.3 demonstrate the importance of treatment cost in the optimal policy. Generally, π_1 is designed to minimize both the visit costs and condition 1 treatment costs. It does this by accepting additional condition 2 treatment costs as a trade-off. As c_1 becomes more expensive, it is intuitive that π_1 would be optimal, as minimizing condition 1 treatments becomes essential, and π_1 does this while also reducing visit costs. As c_2 becomes more expensive, then the trade-off of extra condition 2 visits that π_1 uses to reduce visit costs is no longer worthwhile. If treatment costs are zero, then the decision maker is only concerned with minimizing visit costs. From Corollary 3.4.3, we can see that π_1 is optimal in this case, and therefore is the policy that minimizes visit costs.

Corollary 3.4.1 highlights the importance of considering *both* the full threshold and the joint-start threshold. Looking at just the full threshold would be counter-intuitive, as it suggests that even though π_1 minimizes the treatment costs of condition 1, it is not a good policy when treating condition 1 is expensive. Instead, the full threshold tells us that, for a high c_1 *there are certain states* where π_1 is unlikely to be optimal regardless of c_0 . For example, imagine a patient just received treatment for condition 1 in the previous period, and now condition 2 must be treated in the current period. In this case, π_1 would require that you treat condition 1 again, in back-to-back periods, regardless of how long it might be before condition 1 truly requires treatment. Treating condition 1 in back-to-back periods is very unlikely to be optimal, especially if c_1 is high. Instead, the optimal policy is to not synchronize treatment initially, and instead synchronize treatment later on, when condition 1 is closer to its maximum safe treatment interval. In this way, the gap between \bar{c}_0 and $\bar{\bar{c}}_0$ represents a “start-up” cost of following π_1 . If the patient begins treatment in some state outside of S_{π_1} they may not follow π_1 initially, but once they enter a state in this cycle, π_1 is optimal for the remainder of treatment. If the visit cost is above both \bar{c}_0 and $\bar{\bar{c}}_0$, then the patient is willing to pay the start-up cost of following π_1 . If not, then the patient would rather avoid synchronizing treatment initially, and reconsider synchronization at the next visit.

Corollary 3.4.4 *If the maximum safe treatment intervals are multiples of one another, the joint-start synchronization threshold is zero. That is, if $M_2 = kM_1$ for $k \in \mathbb{N}^+$ then $\bar{\bar{c}}_0 = 0$.*

To prove this, we utilize Lemma 3.4.2 to demonstrate which conditions in Theorem 3.4.2 are relevant to $\bar{\bar{c}}_0$.

Lemma 3.4.2 *If $M_2 = kM_1$ then \bar{c}_0 is fully described by Equation B.7.*

The proof of Lemma 3.4.2 is in Appendix B. Given Lemma 3.4.2, we only need to show that Equation B.7 reduces to $c_0 \geq 0$ in order to show that $\bar{c}_0 = 0$. We start by substituting $M_2 = kM_1$ into Equation B.7:

$$\begin{aligned}
c_0 &\geq -c_1 - c_2 \left(\frac{\delta^{\left(\left\lfloor \frac{M_2 - x_2 - M_1}{M_1} \right\rfloor + 1\right)M_1} - \delta^{\left(\left\lfloor \frac{M_2 - x_2 - M_1 + x_1}{M_1} \right\rfloor + 1\right)M_1 - x_1}}{1 - \delta^{M_1 \left\lfloor \frac{M_2}{M_1} \right\rfloor}} \right) \left(\frac{1 - \delta^{M_1}}{1 - \delta^{M_1 - x_1}} \right) \\
\implies c_0 &\geq -c_1 - c_2 \left(\frac{\delta^{\left(\left\lfloor \frac{(k-1)M_1 - x_2}{M_1} \right\rfloor + 1\right)M_1} - \delta^{\left(\left\lfloor \frac{(k-1)M_1 - x_2 + x_1}{M_1} \right\rfloor + 1\right)M_1 - x_1}}{1 - \delta^{M_1 \left\lfloor \frac{kM_1}{M_1} \right\rfloor}} \right) \left(\frac{1 - \delta^{M_1}}{1 - \delta^{M_1 - x_1}} \right) \\
\implies c_0 &\geq -c_1 - c_2 \left(\frac{\delta^{\left(\left\lfloor \frac{(k-1)M_1 - x_1 - lM_1}{M_1} \right\rfloor + 1\right)M_1} - \delta^{\left(\left\lfloor \frac{(k-1)M_1 - lM_1}{M_1} \right\rfloor + 1\right)M_1 - x_1}}{1 - \delta^{kM_1}} \right) \left(\frac{1 - \delta^{M_1}}{1 - \delta^{M_1 - x_1}} \right).
\end{aligned}$$

We next show that the coefficients on both c_1 and c_2 are negative. The c_1 coefficient is -1 , which is immediate. For the c_2 coefficient, we consider the numerators and denominators of the two fractions individually. Since $\delta \in (0, 1)$, the second numerator and both denominators are clearly positive. Next we note that

$$\begin{aligned}
\left(\left\lfloor \frac{(k-1)M_1 - lM_1}{M_1} \right\rfloor + 1 \right) M_1 - x_1 &= \left(\left\lfloor \frac{(k-1)M_1 - lM_1}{M_1} \right\rfloor + \frac{x_1}{M_1} \right) M_1 \\
&> \left(\left\lfloor \frac{(k-1)M_1 - x_1 - lM_1}{M_1} \right\rfloor + 1 \right) M_1,
\end{aligned}$$

which implies that the first numerator is also positive. Thus, the entire coefficient is negative. Then, since c_1 and c_2 are non-negative by definition, we can reduce this threshold to $c_0 \geq 0$. We have then shown that $\bar{c}_0 = 0$ if $M_2 = kM_1$. \blacksquare

Corollary 3.4.4 is another example of the importance of considering both the full and joint-start threshold values. The high-level result from Corollary 3.4.4 is intuitive. If the maximum safe treatment intervals are multiples of each other, you do not need to adjust the timing of the second condition to synchronize the visits. In a sense they are naturally aligned, and thus there is no downside to synchronizing visits. However, the result of Corollary 3.4.4 only holds for the joint-start threshold, because it relies on the fact that patients within this synchronization cycle do not enter a majority of the possible states (i.e. S_{π_1} is relatively small). However, the full threshold does *not* ignore any states and always applies all of the equations within Theorem 3.4.2. As a result, \bar{c}_0 is almost always non-zero, and can be counter-intuitive if considered in isolation. The

full threshold could be interpreted as saying that even if condition 2 can naturally be synchronized with condition 1, you should not do so. However what a non-zero full threshold really implies is that even if the maximum safe treatment intervals are naturally synchronized, if patients start in a state where the time since last treatment has been staggered (e.g. one of the conditions is offset from this natural synchronization), you may wish to initially deviate from π_1 to align the treatment timing of the conditions for the future. That is, an additional condition 2 visit now might allow for lower costs across the treatment horizon.

3.4.4 Using a heuristic for large state spaces

Solving dynamic programs using policy iteration has a computational complexity of $O(|A||S|^2 + |S|^3)$, where $|A|$ is the size of the action space and $|S|$ is the size of the state space. [75] For this problem, the size of our state space $|S| = \prod_{i=1}^N M_i$. From a practical perspective, we can then see that the computational complexity of this problem can increase dramatically as either the number of conditions N increases or the length of the MSTIs M_i increase.

For conditions such as AMD, new drugs are often coming out which are designed to reduce treatment burden by extending the feasible MSTIs for patients. For example, faricimab is a newer anti-VEGF treatment which has shown potential for extending the maximum MSTI from 12 to 16 weeks. [37] As the scheduling windows grow, we can reduce the burden of computational complexity on clinicians who need to make quick scheduling recommendations for patients during a clinical visit by providing a heuristic for solving the 2-condition problem. In the following section, we outline this heuristic, which is a combination of 3 simple policies that can be solved offline. We then bound the regret of this heuristic and discuss the effect of different problem parameters on this bound.

3.4.4.1 Baseline decision policies

For the 2-condition case, we highlight 3 baseline decision policies that comprise our decision heuristic. The first of these is π_1 , which is defined in Definition 3.4.2. We can anecdotally refer to π_1 as the “cycle 1” policy, as it cyclically synchronizes treatment using visit intervals of length M_1 . The other two policies we refer to as π_2 and π_I , which we define below.

The π_2 policy, is similar to π_1 , in that it reduces visit costs by synchronizing visits. In contrast to π_1 , this policy reduces visits by synchronizing treatment on the frequency of condition 2 (hence the 2 subscript). In summary, π_2 is designed to reduce visit costs while minimizing condition 2 treatment costs. We can anecdotally refer to π_2 as “cycle 2” as it synchronizes treatment on a visit cycle of length M_2 .

Definition 3.4.4 *In the $N = 2$ condition scenario, we define a decision policy $\pi_2(X)$ via the following:*

- *If $X_2 = M_2$ then $a_1 = 1$ and $a_2 = 1$*
- *Else if $X_1 = M_1$ then $a_1 = 1$ and $a_2 = 0$*
- *Else $a_1 = 0$ and $a_2 = 0$*

We next define the π_I policy, which can be anecdotally referred to as an “independent” policy. The π_I policy treats each condition separately with no regard to visit cost. This policy only treats each condition at its MSTI. Generally, π_I is closest to current practice for conditions such as AMD. While π_I does not consider visit costs, it has the advantage of minimizing treatment costs, which makes it a useful policy in specific scenarios (e.g. if treatment is very expensive relative to visit costs).

Definition 3.4.5 *We define a decision policy $\pi_I(X)$ via the following:*

- *For all $i \in [N]$:*
 - *If $X_i = M_i$ then $a_i = 1$*
 - *Else $a_i = 0$*

After defining our baseline policies π_1, π_2 , and π_I , we can now describe our heuristic policy π_H . For each possible state X , the heuristic policy looks at the cost to go $V^\pi(X)$ of the baseline policies. It then selects the action from the baseline policy with the lowest cost to go for that state X . We define this formally below.

Definition 3.4.6 *We define a decision policy π_H such that $\pi_H(X) = \arg \min_{\pi} \{V^\pi(X)\}$ for $\pi \in \{\pi_1, \pi_2, \pi_I\}$.*

3.4.4.2 Bounding the heuristic regret

In this section, we provide a bound on the optimality loss from using the heuristic policy π_H versus solving for the full optimal policy. To do this, we first define a super-optimal cost-to-go value $V^*(X)$. We define $V^*(X)$ such that it is as good or better than the following the optimal decision policy (i.e. it may be unachievable in practice) and such that its value is describable in closed-form. Having a closed-form description of $V^*(X)$ allows us to make a direct comparison to our heuristic value $V^{\pi_H}(X)$. Since $V^*(X)$ is as good or better than the optimal cost-to-go, it follows that the difference $V^*(X) - V^{\pi_H}(X)$ is greater than or equal to the difference between our heuristic and the optimal policy. This difference can therefore serve as a regret bound, which we outline in Theorem 3.4.3.

Theorem 3.4.3 We can define a regret bound for following π_H in the 2-condition joint-start scenario as $R(\pi_H)$, where

$$R(\pi_H) = \min \left\{ c_2 \left(\frac{1}{1 - \delta^{M_1 \lfloor M_2/M_1 \rfloor}} - \frac{1}{1 - \delta^{M_2}} \right), \right. \\ \left. (c_0 + c_1) \left(\frac{\delta^{M_2} - \delta^{M_1 (\lfloor \frac{M_2-1}{M_1} \rfloor + 1)}}{(1 - \delta^{M_1})(1 - \delta^{M_2})} \right), \right. \\ \left. c_0 \left(\frac{1}{1 - \delta^{M_2}} - \frac{1}{1 - \delta^{LCM(M_1, M_2)}} \right) \right\} \quad (3.6)$$

To define a super-optimal cost-to-go $V^*(X)$ we consider the following ideal scenario: a decision policy that fully minimizes both visit costs (like π_1) and fully minimizes treatment costs (like π_I). In general, minimizing both of these costs is unachievable except in specific scenarios. Recall that π_1 minimizes visit costs by requiring extra treatment for condition 2, and therefore by intentionally *not* minimizing treatment costs. Then, assuming that we start in a period where both conditions are treated together (e.g. a “joint-start” patient, who can be described as starting in $X_0 = (M_1, M_2)$), we have the following:

$$V^*(M_1, M_2) = \frac{c_0}{1 - \delta^{M_1}} + \sum_{i=1}^N \frac{c_i}{1 - \delta^{M_i}}.$$

Here, we only require visit costs every M_1 periods, and treatment costs for each condition i every M_i periods. Again, this is generally impossible, as we cannot have a treatment cost in a period without an additional visit cost, unless it aligns naturally.

Our heuristic policy π_H is a combination of the baseline policies π_1 , π_2 , and π_I . Therefore, we know that the cost-to-go of our heuristic policy V^{π_H} cannot be higher than the cost-to-go of any of the baseline policies. That is $V^{\pi_H}(X) \leq \min\{V^{\pi_1}(X), V^{\pi_2}(X), V^{\pi_I}(X)\}$. When $X = (M_1, M_2)$, we can write the cost-to-go values of our baseline policies in closed-form. We have:

$$V^{\pi_1}(M_1, M_2) = \frac{c_0 + c_1}{1 - \delta^{M_1}} + \frac{c_2}{1 - \delta^{M_1 \lfloor M_2/M_1 \rfloor}} \\ V^{\pi_2}(M_1, M_2) = \frac{c_2}{1 - \delta^{M_2}} + \sum_{i=1}^{\lfloor \frac{M_2-1}{M_1} \rfloor} \delta^{M_1 i} \left(\frac{c_0 + c_1}{1 - \delta^{M_2}} \right) \\ V^{\pi_I}(M_1, M_2) = \frac{c_0 + c_1}{1 - \delta^{M_1}} + \frac{c_0 + c_2}{1 - \delta^{M_2}} - \frac{c_0}{1 - \delta^{LCM(M_1, M_2)}}$$

In the π_I cost-to-go, $LCM(M_1, M_2)$ represents the least common multiple of M_1 and M_2 . This

subtraction of visit costs represents visits where both condition 1 and condition 2 naturally align to be treated in the same period, and would otherwise be “double-counted”.

Then, a regret bound $R(\pi_H)$ for a joint-start patient can be written as

$$R(\pi_H) = V^*(M_1, M_2) - \min\{V^{\pi_1}(M_1, M_2), V^{\pi_2}(M_1, M_2), V^{\pi_I}(M_1, M_2)\}$$

Then, by substituting in the values for $V^\pi(M_1, M_2)$ we can write our regret bound as

$$R(\pi_H) = \min \left\{ \begin{aligned} &c_2 \left(\frac{1}{1 - \delta^{M_1 \lfloor M_2/M_1 \rfloor}} - \frac{1}{1 - \delta^{M_2}} \right), \\ &(c_0 + c_1) \left(\frac{\delta^{M_2} - \delta^{M_1 (\lfloor \frac{M_2-1}{M_1} \rfloor + 1)}}{(1 - \delta^{M_1})(1 - \delta^{M_2})} \right), \\ &c_0 \left(\frac{1}{1 - \delta^{M_2}} - \frac{1}{1 - \delta^{LCM(M_1, M_2)}} \right) \end{aligned} \right\}$$

■

From 3.4.3, we can gather insight into the effect of visit and treatment costs on heuristic performance. First, we can see that if the cost of treating condition 2 is 0, the regret also goes to 0. In this case, following π_1 would minimize visit costs with no downside in terms of treatment cost, and would therefore be the optimal policy. Similarly, if the visit cost is 0, the regret goes to 0. In this case, the only goal is to minimize treatment costs. Since the independent policy π_I , does just that, it becomes the optimal policy.

This regret bound also provides interesting insight into the effect of the MSTIs on policy performance. If M_2 gets large while M_1 remains constant, we can see that the regret term associated with π_2 goes to zero. If both M_1 and M_2 get large, from all three of the baseline policy regret values, we can see that the heuristic regret depends less on the absolute size of the MSTIs and more on their relative values respective to one another. For example, the π_I regret term has the $LCM(M_1, M_2)$ factor. Because of this term, MSTIs that align naturally (e.g. 4 weeks and 8 weeks) will have a regret bound of 0. A similar effect can be seen due to the $M_1 \lfloor M_2/M_1 \rfloor$ term in the other individual regret bounds. For this reason, in our numerical analysis in Section 3.5, we describe the effects of an MSTI alignment term, defined as $M_2 \bmod M_1$.

Another advantage of Theorem 3.4.3 is that it provides regret bounds for the individual baseline policies. The components inside of the minimization function represent the regret bounds for π_1, π_2 , and π_I , respectively. In current practice, most clinicians follow π_I and treat conditions independently. The regret bound shown here also allows us to calculate a regret bound for current practice in closed form.

3.5 Numerical Results

In this section we demonstrate the value of our framework through a case study on bilateral AMD. We provide an example patient, and discuss the improvement gained from using the optimal policy versus current practice. We also discuss the performance of the baseline policies. We then study the effect of problem parameters (e.g. visit costs, treatment costs, etc.) on policy performance. We also study the effect of the problem parameters on the synchronization thresholds \bar{c}_0 and $\bar{\bar{c}}_0$. Finally we discuss the potential impact of following our optimization based framework on population level AMD costs.

Throughout our analysis, we study a range of treatment and visit costs. At the Kellogg Eye Center, the most common treatment options available for AMD are bevacizumab and aflibercept. These are anti-VEGF medications designed to prevent the progression of AMD. Individual patients respond differently to each medication, with some medications being ineffective for some patients. Very recently, a newer drug called faricimab has also entered the market, with potential of enabling longer MSTIs for patients. [37] Per our clinical collaborators, the cost of a single injection for each drug at the Kellogg Eye Center is \$50 for bevacizumab, \$1,300 for aflibercept, and \$1,800 for faricimab. The dramatic price discrepancy comes from the fact that bevacizumab was not originally designed for AMD treatment (in contrast to the other two medications). However, clinical trials have shown bevacizumab to be as effective as the other medications despite its off-label use. [31, 10, 13] Consequently, clinicians typically prescribe bevacizumab to patients as a first attempt to prevent disease progression. If patients do not respond to this drug, clinicians then use a step-therapy approach, trying more expensive drugs as needed. In our analysis, we test the treatment cost values in the \$0-\$1,800 range. Note that a \$0 treatment cost might represent a patient who has full insurance coverage for the injection. Visit costs at the Kellogg Eye Center are currently billed at \$244 per visit. Note that this cost only includes administrative fees, and does not include personalized patient costs, such as travel or caregiver costs. As a result, we test a range of patient visit costs from \$0-\$1,000.

Other problem parameters include the patient's MSTIs. The MSTIs for individual eyes diagnosed with neovascular AMD currently range from 4 to 12 weeks, and are treated using weekly intervals. We study all weekly scheduling options within this range. While the specific MSTI values are important parameters, an additional useful consideration is how well the MSTI values align. For example, MSTIs of 4 and 8 weeks are quite easy to schedule, as they naturally align. In contrast, the best policy for patients with MSTIs of 4 and 11 weeks is much less intuitive. As a result, we study the effect of MSTI alignment on policy performance using the quantity $M_2 \bmod M_1$ as a measure of the level of alignment.

The appropriate discount factor δ for a patient is difficult to quantify, so we test a range of dis-

count factors between 0.01 and 0.99. For individual patients, we turn to our clinical collaborators for guidance. Generally, the future is of high value, however since treating within the MSTI does not cause disease progression, the future costs are mostly monetary (versus health-based). As a result, maintaining a relatively high discount factor (such as $\delta \in (0.990, 0.999)$ used in Chapter 2) is considered less important in this analysis.

3.5.1 Example patient

Consider a patient with bilateral AMD who has MSTIs of 4 weeks and 7 weeks. Recall that bevacizumab is generally the first medication administered to a patient due to its relatively low cost and high efficacy. [13] Assume that the patient responds to bevacizumab treatment and pays \$50 per injection (i.e. $c_1 = c_2 = \$50$). Additionally, assume that the patient only pays the \$244 administrative visit fee, and does not consider other visit costs such as travel or the value of time. For a joint-start patient, we can see the outcome of following different treatment policies in Table 3.1 below. Table 3.1 shows us that following an optimal decision policy could reduce patient costs by \$2,064 (in terms of discounted lifetime costs) versus the current practice policy of π_I . Additionally, we can see that for this patient, the cost of following optimal policy is equivalent to both π_1 and the heuristic policy π_H .

Policy	π_*	π_1	π_2	π_I	π_H
Discounted Lifetime Cost	\$8,730	\$8,730	\$9,220	\$10,794	\$8,730

Table 3.1: The discounted cost of following different decision policies for a joint-start bilateral AMD patient being treated with bevacizumab.

For this same example patient, we can also calculate the synchronization thresholds \bar{c}_0 and $\bar{\bar{c}}_0$ as \$300 and \$103, respectively. Recall that the more practical threshold for interpretation is $\bar{\bar{c}}_0$. We see that this patient has a visit cost above $\bar{\bar{c}}_0$, which is why the π_1 policy is optimal for this joint start patient. Interestingly, the patient's \$244 visit cost is not above the \bar{c}_0 threshold. This means that if the patient had a different initial state, they may not have followed π_1 exactly as the optimal policy. For examples such as this, where the patient's visit cost is in between the thresholds, the optimal policy only differs from π_1 at the start of treatment. Once the patient reaches a state inside of S_{π_1} , they follow π_1 for the remainder of treatment.

3.5.2 Heuristic and baseline policy analysis

In this section, we show how varying each of the problem parameters affects the average policy performance for the optimal, heuristic, and baseline policies. Figure 3.4 below shows the average

performance effect across a range of visit costs, treatment costs, and discount factors for a joint-start patient. Additionally, Figure 3.4 shows the effect of the alignment between the MSTIs, as measured by $M_2 \bmod M_1$.

Across all of the graphs in Figure 3.4, we see that our heuristic policy is very robust to the problem parameters, and performs nearly optimal across all test cases. In our simulation, we tested 20,520 parameter combinations, and the heuristic was equivalent to the optimal policy in 19,887 (96.9%) of those combinations. For the combinations where π_H was sub-optimal, the maximum regret was 4.0%.

In Figure 3.4, we see that as the visit cost increases, the π_1 policy gets closer to the optimal policy. This matches the intuition derived from Theorem 3.4.2. We also see that for small visit costs, the independent policy π_I gets closer to optimal. This is intuitive, because for low visit costs, minimizing treatment is most important, and π_I minimizes treatment costs. In fact, when the visit cost is \$0, π_I is optimal.

As the treatment cost for condition 1 increases we see that both the π_1 and π_I policies perform closer to optimal. The π_2 policy starts to perform worse, because π_2 reduces visit costs by scheduling additional condition 1 treatment. As condition 1 treatment becomes more expensive, this trade-off becomes less valuable. We see similar results regarding the treatment cost of condition 2. In this case, π_1 starts to perform worse as the cost increases, again due to the trade off of visit and treatment costs for condition 2. Notably, if condition 2 treatment costs are \$0, then π_1 is optimal. This makes sense, because π_1 minimizes both visit costs and condition 1 treatment costs, which would be the only cost components in this scenario.

As the discount factor approaches 1, the regret for each of the baseline policies grows dramatically. This is primarily because repeated deviations from the optimal policy become more important as the value of the future increases—therefore even slight policy differences start to have a large impact on cost. Notably, our heuristic policy remains robust even as the discount factor gets very close to 1. The effect of the MSTI alignment depends on the policy of interest. If $M_2 \bmod M_1$ is small, this means that shortening the treatment interval of the second condition (which π_1 does) does not have as large of an impact on cost. Therefore, for a smaller alignment metric, π_1 performs better. In contrast, if the MSTI alignment metric is large, this means that adding in an extra condition 1 visit to synchronize with condition 2 (which π_2 does) happens less frequently and therefore has a smaller cost impact. Therefore, for larger $M_2 \bmod M_1$, we see that π_2 starts to perform better. Again, our heuristic balances the best of both of these policies and performed well for all of the alignment metrics in our analysis.

In Figure 3.5, we analyze the contribution of each of the baseline policies to our regret bound for a joint start patient. From Equation 3.6, we see that the actual regret bound is the minimum of the 3 lines shown in each of the plots within Figure 3.5. Similar to the results found in Figure 3.4,

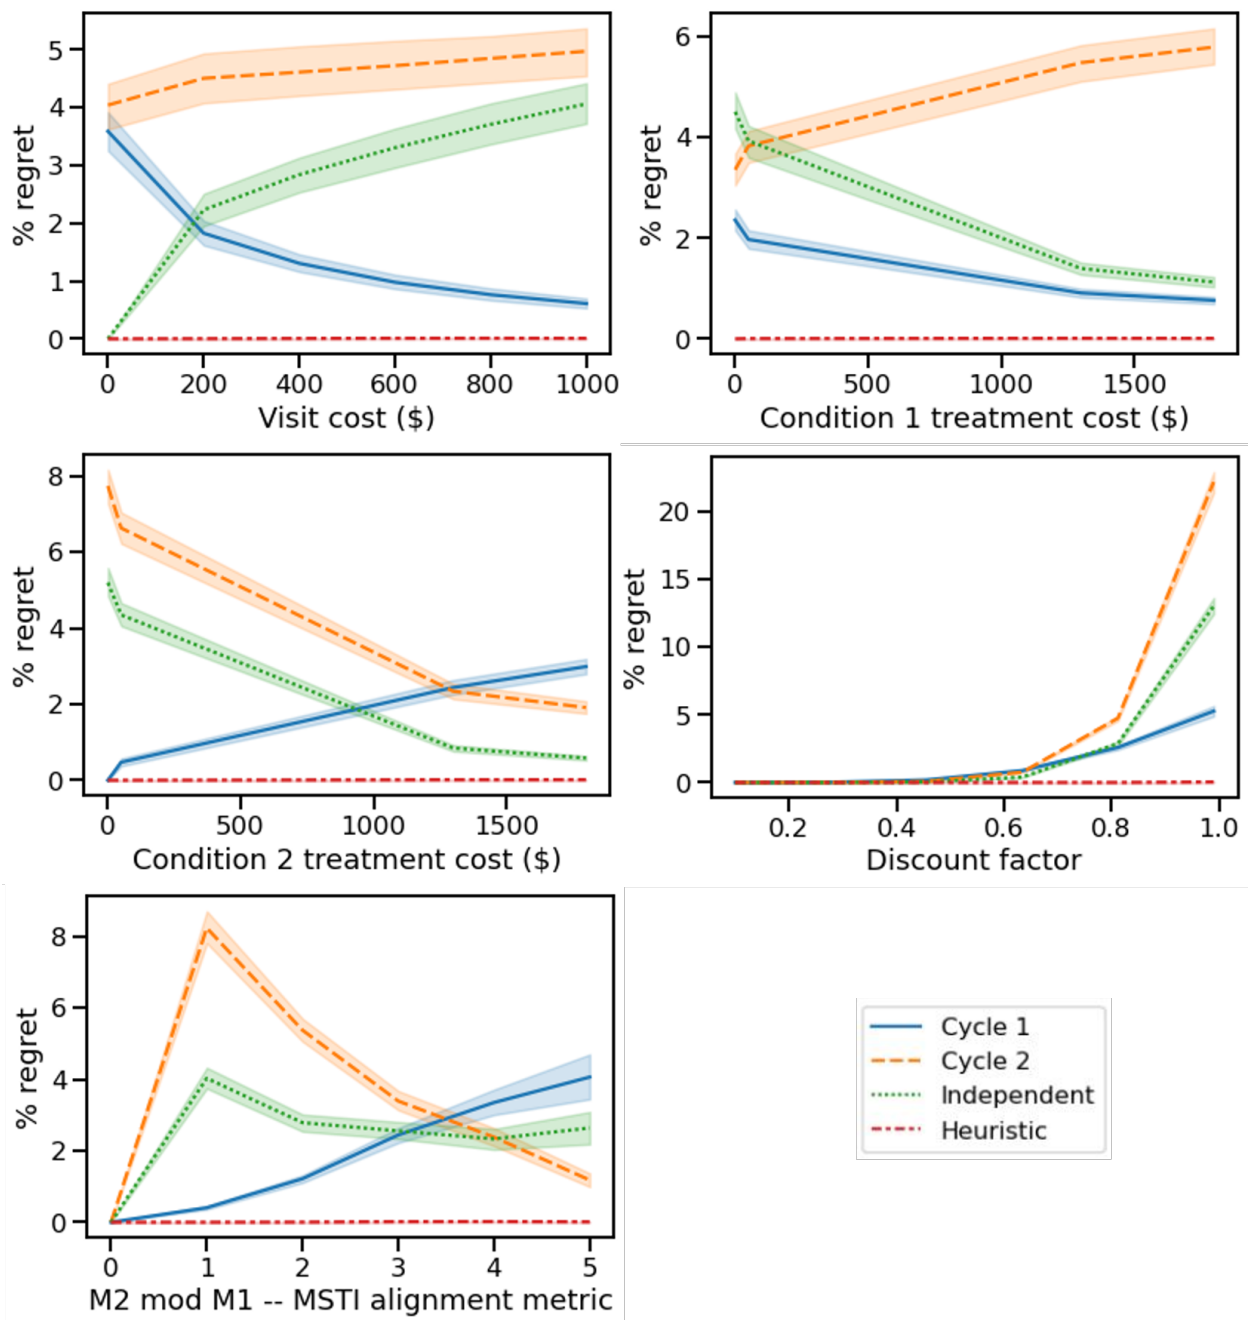


Figure 3.4: Plots comparing the performance of the baseline decision policies versus the optimal policy. Performance is measured in terms of regret (i.e. the percent increase in cost compared to the optimal decision policy).

we see that a discount factor close to one has arguably the largest effect on regret, and therefore the regret bound as well. However, even in this scenario our measured regret was only as high as 4.0%.

3.5.3 Synchronization threshold analysis

In this section, we analyze the effect of the patient-specific parameters on the synchronization thresholds \bar{c}_0 and $\bar{\bar{c}}_0$, and highlight the key differences between the thresholds.

Figure 3.6 shows both the full and joint-start thresholds for an AMD patient receiving bevacizumab (i.e. $c_1 = c_2 = \$50$) with a discount factor of 0.99. From this figure, we can see the results of Corollary 3.4.4, which say that when $M_1 = M_2$, the joint-start threshold is zero. Interestingly, we see that for AMD patients, the majority of MSTI combinations have a joint-start threshold of 0. This suggests that for many patients, the π_1 policy is always optimal, regardless of the visit cost. The combinations where π_1 is not optimal are generally those where reducing the treatment interval of condition 2 to match condition 1 would have a dramatic increase in treatment frequency. For example, treating a 7 week condition every 4 weeks would nearly double the treatment frequency, and is less likely to be optimal. This matches the results in Figure 3.4, which suggest that a high $M_2 \bmod M_1$ was associated with poor π_1 performance. Additionally, Figure 3.6, highlights a dramatic difference in the full and joint-start thresholds. This demonstrates the practical importance of the joint-start threshold when deriving policy insights.

In Figure 3.7 we show the effect of treatment costs and the discount factor on the synchronization thresholds. Because the thresholds depend significantly on the patient’s MSTI, we consider an example patient with MSTIs of 4 and 7 weeks to isolate the effect of treatment cost and the discount factor. Additionally we only display a range of discount factors from $\delta \in (0.50, 0.99)$. For discount factors very close to 0, the thresholds get extremely large, and make visual interpretation of the parameter effect very difficult. We can still derive a practical insight from this effect—for very small δ , the future is considered unimportant. Therefore, the optimal policy is a greedy policy, which only focuses on the current period. In that case, you should only treat a condition if required in this period, which makes π_1 almost never optimal.

From Figure 3.7, we can see that our numerical analyses confirm our findings from Corollaries 3.4.1 and 3.4.2. As the cost of treating condition 1 goes up, π_1 is more likely to be optimal for joint-start patients. In this analysis, we see that for AMD patients, the threshold goes to zero quite quickly, even for poorly aligned MSTIs like the 4- and 7-week combination in our example. Additionally, we see that as the cost to treat condition 2 goes up, π_1 is almost never optimal, and instead the patient should follow a policy closer to π_2 .

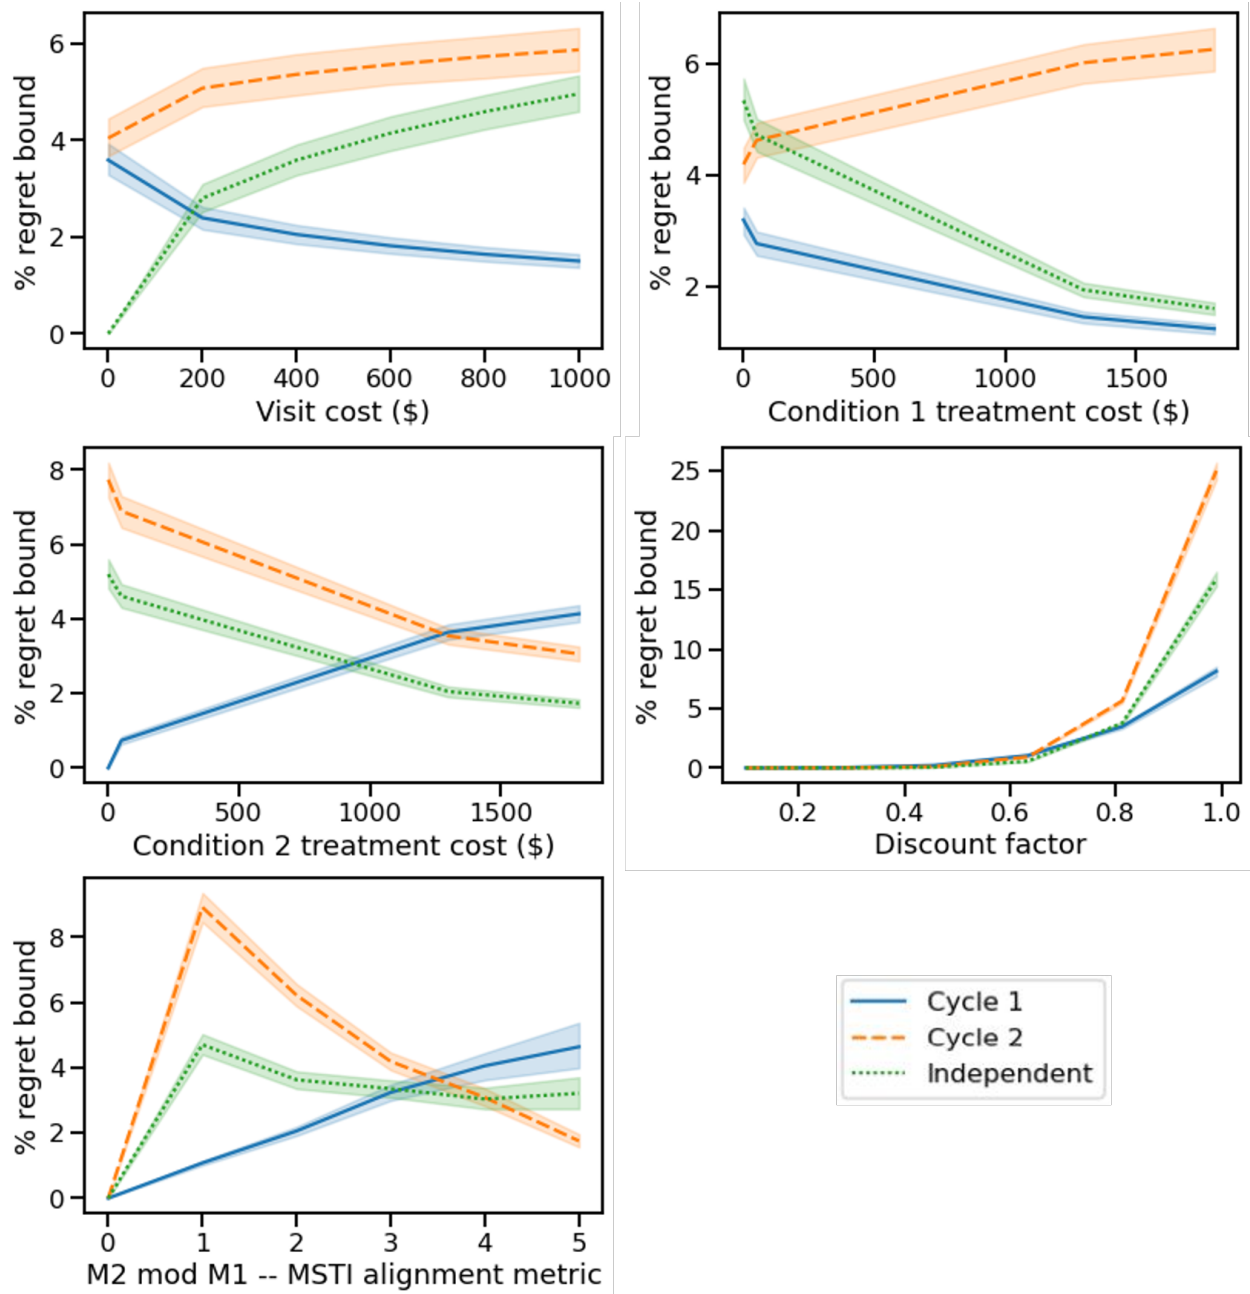


Figure 3.5: Plots comparing the performance of the baseline decision policies versus the optimal policy. Performance is measured in terms of regret (i.e. the percent increase in cost compared to the optimal decision policy).

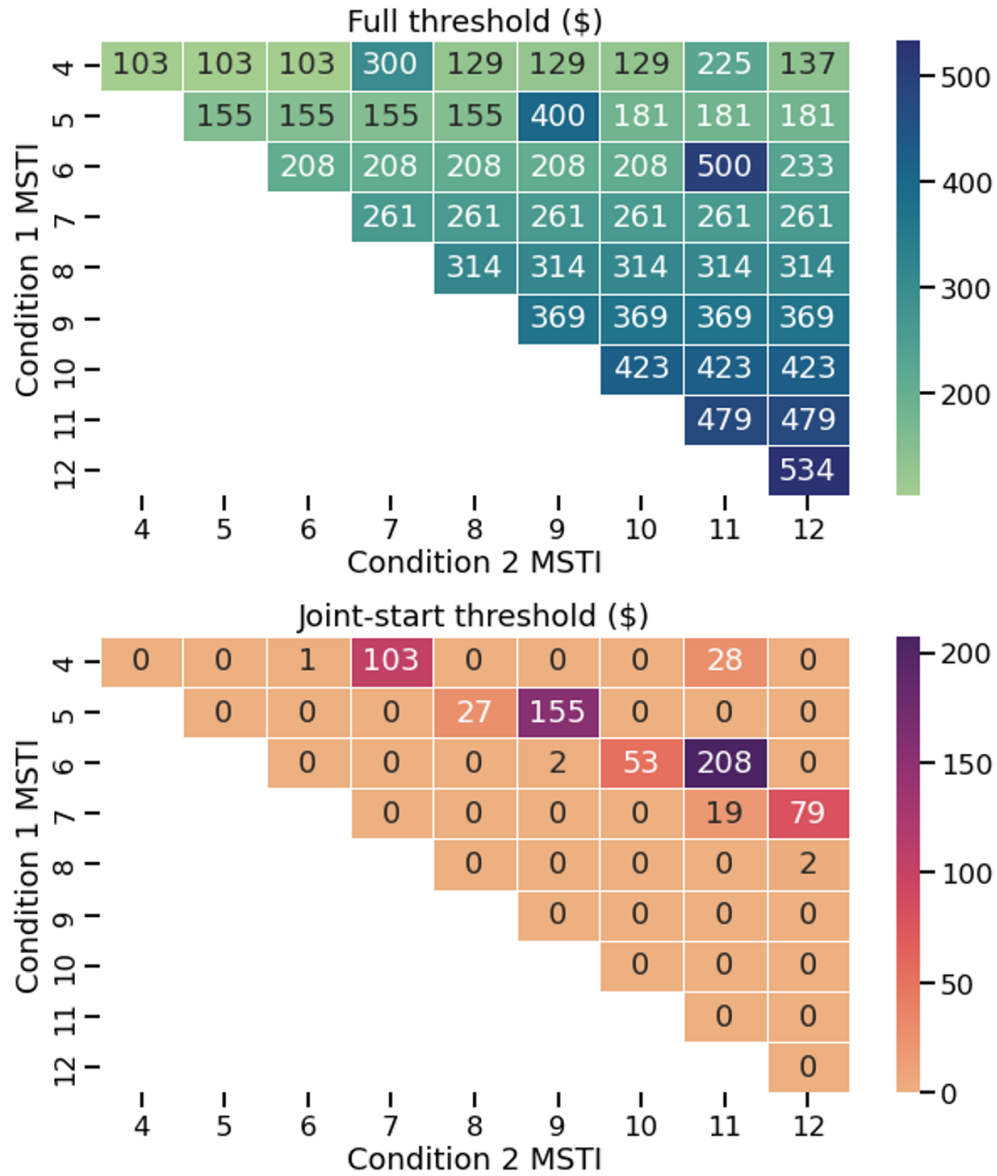


Figure 3.6: The full (\bar{c}_0) and joint-start ($\bar{\bar{c}}_0$) synchronization thresholds across a range of maximum safe treatment intervals for an AMD patient receiving bevacizumab.

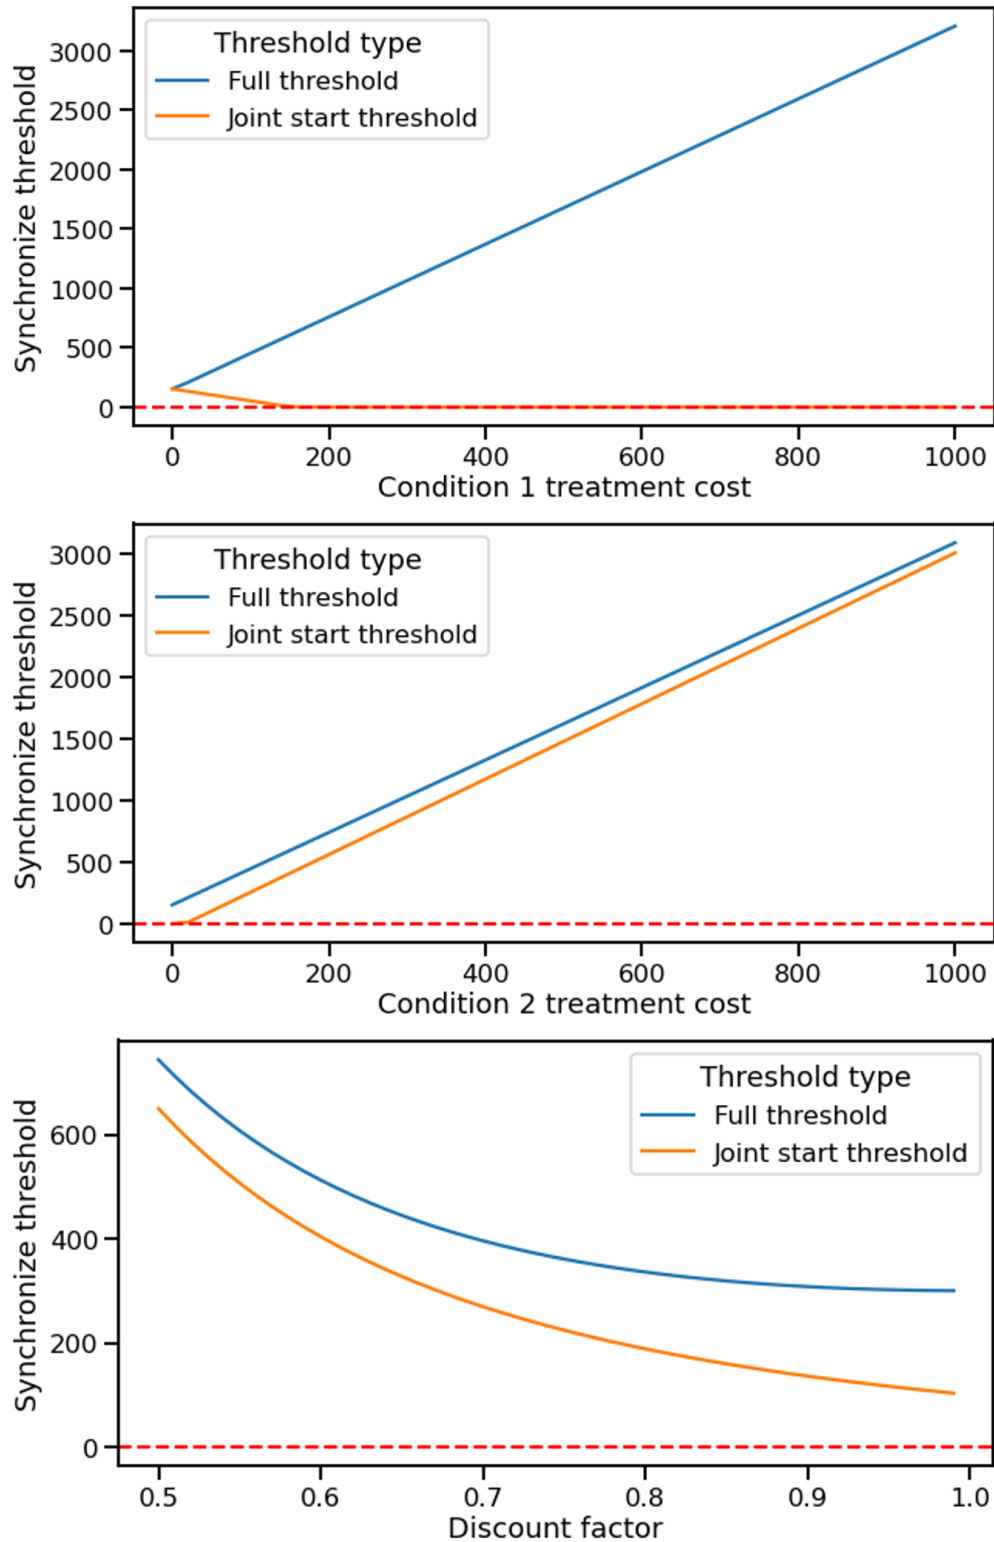


Figure 3.7: The full (\bar{c}_0) and joint-start ($\bar{\bar{c}}_0$) synchronization thresholds for an AMD patient with MSTIs of 4 and 7 weeks. For each plot, we assume baseline treatment costs of $c_1 = c_2 = \$50$, and a discount factor of $\delta = 0$, and vary only the parameter of interest.




Figure 3.8: The distribution of MSTI combinations for patients in the Kellogg dataset. Blank spaces indicate that no patients had that MSTI combination.

3.5.4 Population level analysis

In this section, we use patient data from the Kellogg Eye Center at the University of Michigan to estimate the potential population level benefit of following our treatment scheduling framework. For our analysis, we looked at patients with bilateral AMD being treated with either bevacizumab or aflibercept. We only considered patients with known MSTIs in both eyes, and whose MSTIs were not equal. We only considered unequal MSTI combinations because Corollary 3.4.4 shows that the scheduling problem is trivial if the MSTIs are equal. In our dataset, 186 patients met these criteria, and the distribution of the MSTIs are shown in Figure 3.8.

Using each individual patient’s MSTI values and treatment costs, we calculated the lifetime cost of following each of the baseline decision policies. Because we did not have access to individual visit cost information, we used the baseline visit cost of \$244 per visit, which only includes the administrative cost of a visit, and ignores external factors, such as travel and caregiver costs. Additionally, we used a discount factor of $\delta = 0.99$ based on clinical recommendation.

Figure 3.9 shows the discounted lifetime cost distributions of following each policy for both bevacizumab and aflibercept patients. Recall that current practice is to schedule treatment using the independent policy π_I . In Figure 3.9, we see that for patients receiving bevacizumab injections, following the optimal policy offers a significant improvement over current practice. On average, patients saw a cost reduction of \$2,127. Additionally, the optimal policy was equivalent

to both the π_1 and π_H policies. For patients receiving aflibercept injections, the optimal policy was equivalent to the independent treatment policy for almost all patients, with the optimal policy offering an average improvement of only \$50 per patient. This is because aflibercept injections have a significantly higher treatment cost of \$1,300 per injection, versus the \$50 per injection cost of bevacizumab. In this scenario, minimizing treatment costs is extremely important. Across all patients and medications, the average improvement in discounted lifetime cost was \$1,189 over current practice.

An estimated 3 million U.S. adults have advanced AMD with a third of these patients diagnosed with bilateral AMD. [13] In the Kellogg dataset, 49% of patients had different MSTIs between eyes. Extrapolating the results of the Kellogg patient data to a population level, these results suggest that following an optimization based approach to treatment scheduling could offer a discounted savings of approximately \$582 million in direct medical costs alone. These are estimated lifetime savings for patients currently diagnosed with bilateral AMD. As new diagnoses occur, the potential savings from considering synchronization would continue to grow. Additionally, incorporating indirect costs like transportation and caregiver costs would further increase this estimate.

3.6 Discussion

In this chapter, we introduce a dynamic programming approach to scheduling treatment for patients with multiple chronic conditions. We discuss when it is beneficial for patients to synchronize the treatment of their different conditions, thereby reducing their overall visit costs. We show that for a high enough visit cost, the optimal policy is to always synchronize visits using the π_1 policy, regardless of the number of conditions. For patients with two conditions, such as those with bilateral AMD, we calculate these visit cost thresholds directly. We also analyze how individual patient parameters such as treatment costs and MSTIs affect these thresholds.

One particularly interesting finding of this work is the important distinction between the full and joint-start thresholds \bar{c}_0 and $\bar{\bar{c}}_0$. From a fully theoretical point of view, we cannot call the π_1 synchronized policy optimal unless a patient’s visit cost is above the \bar{c}_0 threshold. However, this threshold is much less useful from a practical point of view. Even following a fully independent scheduling policy, patients will eventually synchronize their visits at some point in the future, reaching a “joint-start” state. After this point, the $\bar{\bar{c}}_0$ threshold is the only meaningful threshold. Consequently, following the traditional definition of an optimal policy can actually be misleading for patients, causing them to over-complicate their scheduling approach. By distinguishing between the two thresholds, we are able to highlight that the optimal policy for many patients is in fact a simple policy, such as π_1 , π_2 or π_I . We also show that an easy to solve heuristic based on these simple baseline policies performed well across a range of problem parameters.

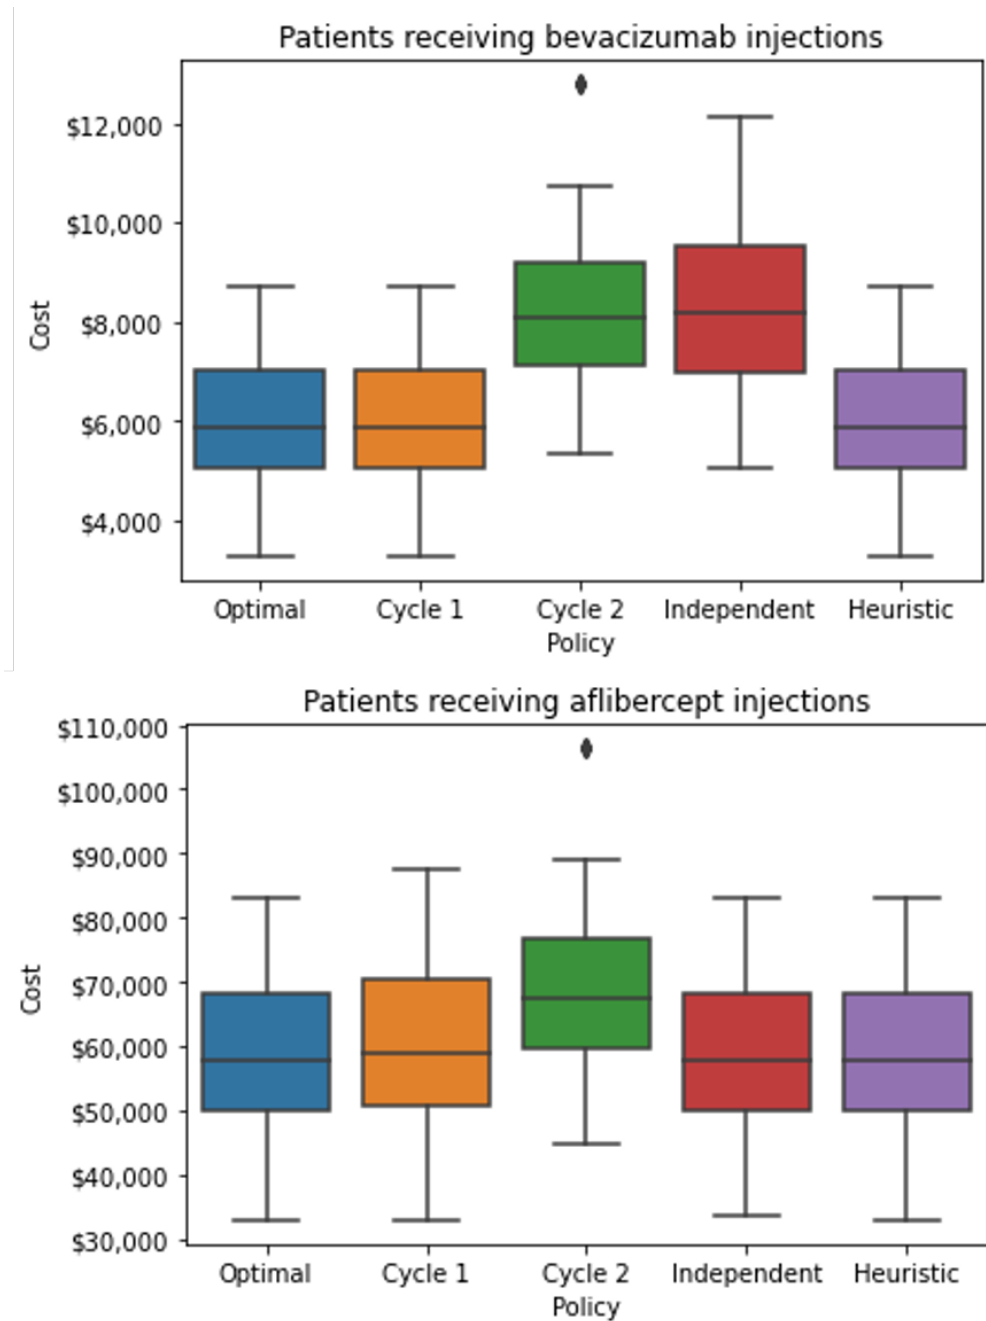


Figure 3.9: The discounted lifetime costs of following each baseline policy across medications.

As an application to AMD, we demonstrate the power of our optimization approach in reducing direct medical costs for patients. From a holistic perspective, we believe that our estimated \$582 million savings is conservative, as it does not include the many indirect or intangible costs that patients with bilateral AMD face. As just one example, AMD is a disease that causes blindness, and for many patients even getting to the clinic can be a massive emotional and financial burden. They often rely on caregivers to transport them to the clinic, or have to miss time from work. Saving even a single visit can have a major impact on a patient's quality of life, not just on their wallet.

Beyond AMD, this work is one of the first approaches to managing multiple chronic conditions. We note that bilateral AMD is one of the best candidates for synchronizing treatment for a number of reasons. For example, both eyes are managed by the same clinician, which means that coordinating the schedules of multiple clinicians is not required to synchronize treatment. Additionally, having the same treatment for both eyes means that there are no potential interaction effects between multiple medications. However, we happily note that our model is able to adapt to scenarios without these features. As one example, if certain medications are incompatible for synchronization (perhaps due to interaction effects), we could limit our available action set A to not allow for combinations that could have a negative impact on patient outcomes. As another example, if conditions managed by different providers are treated in the same period, we could add a coordination fee to the visit cost. Our baseline modeling approach is flexible enough to incorporate a range of scenarios while still providing insights into how best to manage multiple conditions.

Recall that 4 in 10 U.S. adults suffer from multiple chronic conditions. [16] We hope this work is a stepping stone to a modeling paradigm where considering all of a patient's conditions is the norm, not the exception.

CHAPTER 4

Increasing Organ Donation Rates

4.1 Introduction

Organ transplantation is a life-saving and cost-effective intervention for patients for patients with organ failure. The need for organ transplantation far outpaces the rate of organ donation in the US, resulting in over 7,000 waitlisted candidates dying annually. [40] Even patients who ultimately undergo transplantation experience prolonged morbidity due to long wait times. Wait times vary drastically by organ procurement organization (OPO)—for example, the proportion of patients receiving kidney transplants within the first 5 years of listing varies from 10% to nearly 80% across donation service areas. [34]

In this chapter, we analyze two policies with the potential to reduce waitlist mortality by increasing organ donation rates. In Section 4.2, we discuss a presumed consent donation policy in which willingness to donate is the default option. Presumed consent has been implemented internationally, and has been associated with significant increases in donation. [77] However, the effect of an associated increase might have on key patient outcomes is unclear. We simulate the effect this policy would have on waitlist size and mortality across a range of donation rate increases.

In Section 4.3, we discuss the use of “ineligible” donors. These are donors who have already expressed a willingness to donate, but due to specific health criteria are not formally considered good candidates for organ donation. However, the health criteria describing an “eligible” donor are not true requirements for donation. For example, one eligibility criterion is that an organ donor be younger than 76 years old. It is a common practice that transplant physicians accept donations from healthy individuals older than this age if they believe the transplant will have a positive impact on a waiting patient. [40] We study the difference in survival for recipients of eligible versus ineligible donations, and estimate the impact of standardizing ineligible donor use nationally.

4.2 Increasing Donation via Presumed Consent

One potential, albeit controversial, policy that has been proposed to increase organ donation is presumed consent, or an opt-out policy. Such a policy would make willingness to donate the default option, unless an individual explicitly opts not to be an organ donor. [47] Variations of this policy have been adopted in several countries worldwide, with mixed results. [1, 62, 82] Nonetheless, there is evidence that lack of consent plays a role in preventing donation in up to 20-40% of otherwise eligible deceased donors in the US. [88, 29] This, however, is likely an underestimation of the donors missed due to lack of authorization, because it is calculated based on a limited pool of donors pre-defined as having an “eligible death,” and many organ transplants occur from donors outside of this definition. [55] Thus presumed consent has been proposed as a potential avenue to increase organs available for transplant. In this chapter, we develop a model to examine the impact of such a policy on a historical cohort of waitlisted patients to give realistic estimates of the impact of a presumed consent policy in the United States.

4.2.1 Modeling approach

We analyzed data from the Organ Procurement and Transplantation Network (OPTN) from 2004-2014 to develop a waitlist model for solid organ transplant. The data was restricted to adult (18+) solid organ donation candidates (heart, kidney, liver, lung, pancreas), and included information on candidate age, organ(s) requested, date added to the waitlist, date removed from the waitlist, and removal reason. Removal reasons were classified into the categories of 1) received living donor transplant, 2) received deceased donor transplant, 3) died while waiting for transplant or became too ill to receive transplant, and 4) other.

We simulated the impact of a presumed consent policy on waitlist outcomes by organ type in monthly intervals during the 10-year study period. The key metrics used to assess this impact were the number of patients on the waitlist and the number of patient removals due to illness or death. The model simulated the effects of presumed consent by increasing the number of deceased donors that donated at least one organ. For analyses considering all organs, candidates waiting for multiple organs were listed only once. For analyses scoped to individual organs, candidates waiting for multiple organs were listed on all corresponding waitlists. In all instances, candidates who were simultaneously listed at multiple centers for the same organ were only counted once. As new organs became available, the model considered two allocation policies when adjusting waitlist removal rates. Candidates who were listed, underwent transplant, and later relisted were counted multiple times. A schematic of the model can be seen in Figure 4.1.

A systematic review by Rithalia, et al. examining the impact of the institution of presumed consent policy from several countries worldwide, showed that the policy is associated with up to

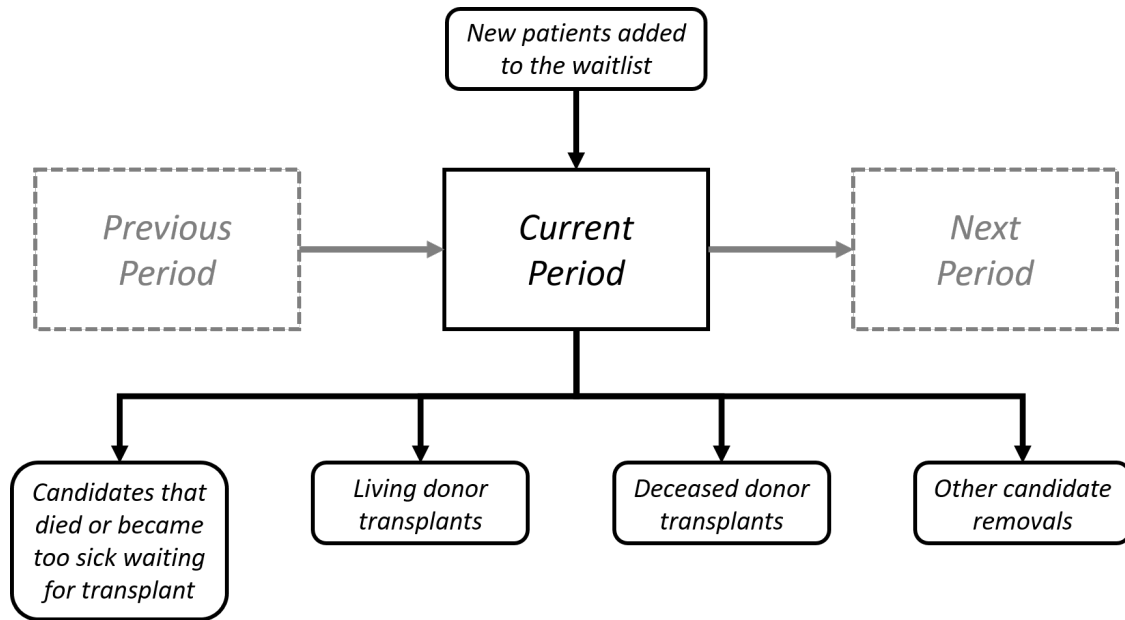


Figure 4.1: An overview of the waitlist model dynamics.

a 25% increase in the deceased donor availability rate. [77] However, more recent studies have reported more modest improvement in deceased donor availability. [82, 67, 99] Accordingly, we modeled a practical range of potential deceased donor increases from 5% to 25% with such a policy change in the US. We refer to a 5% increase in deceased organ donations as the base case for a presumed consent impact. In all cases, the number of living donor donations and waitlist addition quantities were not changed from true historical values.

In order to translate a deceased donor increase to an increase in the number of transplants, we also incorporated organ yield rates into our model. According to the OPTN Annual Data Report, the average annual organ yield for deceased donors was 0.29 hearts per donor, 1.47 kidneys per donor, 0.76 livers per donor, 0.34 lungs per donor, and 0.15 pancreata per donor over the model time period. [41]

When simulating the impact of additional donors, there is significant variation in the allocation of organs to patients. [28] Deciding which individual will receive a newly available organ depends on several factors such as blood type, geography, physician decision making, patient availability, etc. Limitations in our data set and modeling methodology did not allow us to utilize a single prescriptive allocation method that appropriately matched clinical practice. As a result, the study developed two allocation policies which are intended to represent the range of potential outcomes.

In the first allocation policy—referred to as a random policy—additional organs were assigned to currently waiting candidates regardless of patient demographics, such as age, time on the waitlist, medical illness, or waitlist priority. As current clinical allocation policies are designed to help

patients with the most need, a random allocation policy would serve as a lower bound for the possible impact of presumed consent.

In the second allocation policy—referred to as an ideal policy—the model matched additionally available organs with candidates that would have otherwise died or become too sick for transplant in that same time period. Waitlisted patients who were delisted due to death or illness were given newly available organs first, just prior to their waitlist removal. This scenario was designed to emulate the case in which clinicians are able to perfectly forecast patient health and need, and that all additionally available organs are compatible with the sickest patients. This is considered the ideal scenario in terms of impact on patient health and, while unlikely, can serve as an upper bound for the possible impact of presumed consent. The true effect of a presumed consent policy would lie somewhere in between a random and ideal allocation model.

The model output included the estimated number of waitlist candidates in each time period, as well as the number of candidates in each removal category for each organ. This output was combined with published results on organ transplant survival benefit to estimate the potential effect of presumed consent on patient life years gained from transplant. [76]

4.2.2 Numerical results

4.2.2.1 Historical data

From 2004-2014, there were in total 524,359 unique potential organ recipients on the waitlist. The number of candidates on the waitlist on a year to year basis saw wide fluctuation based on the organ type, from an 83% increase in kidney transplant candidates, to a 56% decrease in lung transplant candidates. The number of patients on the waitlist at the beginning of each year is shown in Table 4.1.

Table 4.1: The size of the organ waitlists on January 1st of each respective year.

Year	Heart	Kidney	Liver	Lung	Pancreas	All Organs
2004	3,252	54,779	15,249	3,641	3,665	78,011
2005	2,938	58,507	15,409	3,619	3,777	81,623
2006	2,742	63,010	15,447	2,915	3,868	85,254
2007	2,593	67,759	15,490	2,649	3,805	89,612
2008	2,393	72,946	15,331	2,007	3,643	93,677
2009	2,498	76,660	14,919	1,881	3,565	96,828
2010	2,792	81,801	14,967	1,755	3,424	102,038
2011	2,886	86,307	15,301	1,706	3,384	106,729
2012	2,829	89,399	15,295	1,606	3,173	109,514
2013	3,073	93,526	15,077	1,588	3,069	113,449
2014	3,427	97,774	15,106	1,563	2,968	117,914

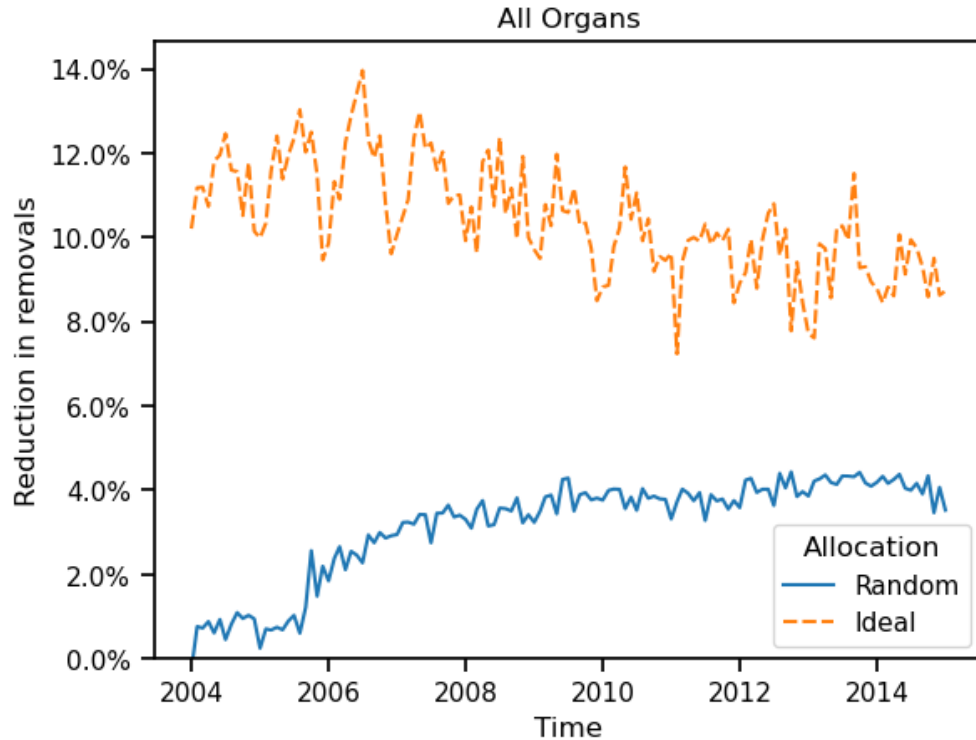


Figure 4.2: Reduction in the number of removals from the waitlist due to death or illness across all organs from 2004-2014 with a 5% increase in donors with presumed consent, stratified by allocation.

4.2.2.2 Presumed consent impact

To determine the impact of presumed consent, we considered a range of increases to the number of deceased donor transplants. Figure 4.2 presents the potential impact of the presumed consent base case (5% increase in donation) on the number of candidates removed due to death or illness under both allocation policies. The impact on the number of removals due to sickness or death varied between scenarios. In the base case, the random allocation policy reduced these removals by an average of 3% over the time period, while the ideal allocation policy would have reduced these removals by an average of 10%.

The ideal allocation policy would have an immediate impact on the number of removals due to death or illness (9%-14%). This impact trended lower over time (8%-10%) because the increase in donors is not enough to completely eliminate removals due to patient death or illness, so the waitlist continued to grow over the time period studied. Conversely, the random allocation policy would have a smaller impact on removals due to death or illness (~1%) that grew to approximately 4% over the time period studied. This is because with a random allocation, the rate of deaths and illness is proportional to the size of the waitlist and as the random policy reduces the size of the

waitlist, it also reduces the number of removals due to deaths or illness. A breakdown of the impact over time on the number of removals due to illness or death by organ type and allocation policy is presented in Figure 4.3 and Figure 4.4. The variation in the curves reflect the variation in monthly removals from the waitlist.

Table 4.2 shows the average monthly reduction in these removals for each organ over a range of presumed-consent related donation increases. Even an ideal allocation scenario with a 25% increase in presumed consent-associated organ donors would not have provided enough donors to eliminate deaths or removals due to illness while waiting for an organ transplant.

Table 4.2: Mean monthly reduction in removals due to sickness or death by organ, 2004-2014

Organ	5% Impact		15% Impact		25% Impact	
	Random Allocation	Ideal Allocation	Random Allocation	Ideal Allocation	Random Allocation	Ideal Allocation
Heart	2%	6%	7%	18%	11%	29%
Kidney	4%	12%	10%	37%	17%	62%
Liver	3%	8%	9%	25%	16%	41%
Lung	4%	7%	10%	22%	16%	36%
Pancreas	1%	3%	2%	10%	4%	16%
Combined	3%	10%	10%	31%	16%	52%

Figure 4.5 shows the number of candidates on the waitlist by allocation policy for the base case. Under an ideal allocation policy, presumed consent would need to generate enough additional organs to eliminate removals due to deaths and sickness before other candidates that would remain on the waitlist are affected. The impact on a random allocation policy on the organ waitlist is shown in Table 4.3. The random allocation policy in the base case scenario did result in a reduction in waitlist quantities, however this is because this policy does not perfectly assign additional organs to patients that would have died or become too sick. The random allocation scenario resulted in a 4%, 13%, and 22% reduction in the number of candidates awaiting solid organ transplant at the end of the study period given a 5%, 15% and 25% impact, respectively. The model using a random allocation policy reduced the combined waitlist growth from 54% to 48% in the base case. The simulation results found that, under the random allocation policy, the impact required to eliminate organ waitlists in the study period varied between organs, ranging from a 97% increase for the lung waitlist to a 373% for the pancreas waitlist. There would be no change in the waitlist with ideal allocation policies because all excess organs would be allocated to patients who would otherwise left the waitlist, even in a 25% increase scenario.

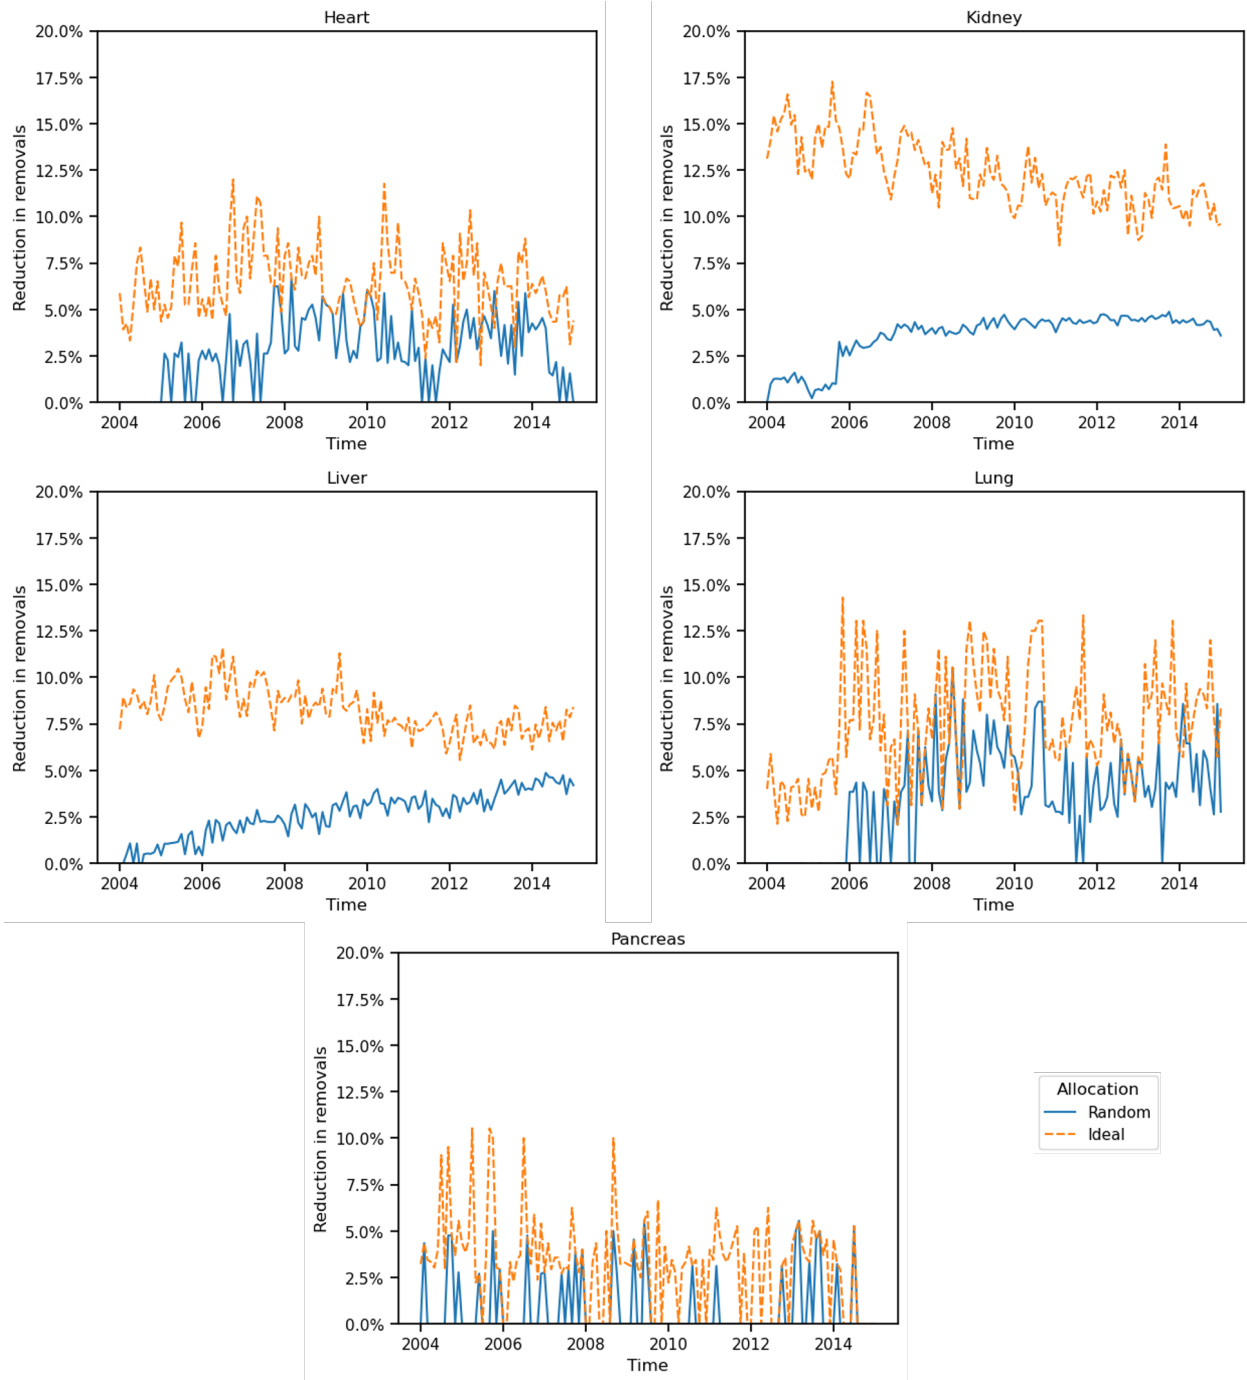


Figure 4.3: Removals from the waitlist due to death or illness from 2004-2014 with a 5% increase in donors with presumed consent, stratified by organ and allocation.




Figure 4.4: Reduction in the number of waitlist removals due to death or illness by organ

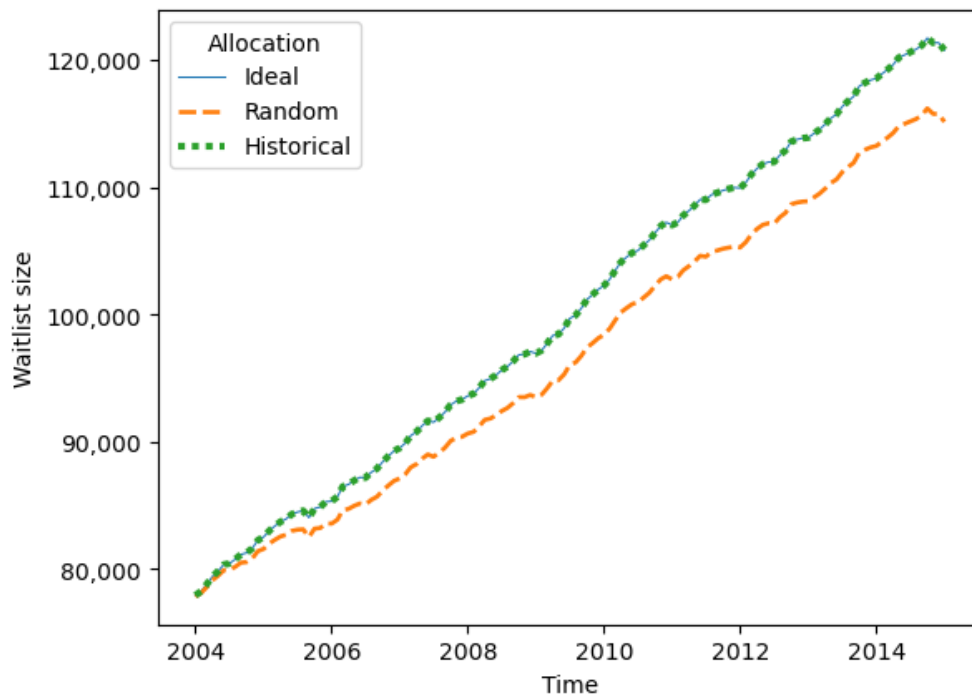


Figure 4.5: Overall number of waitlist candidates stratified by allocation.

Table 4.3: Percent change in the number of waitlist candidates from 2004-2014 under a random allocation policy

Organ	Opt-In Policy (Historic)	Presumed Consent Impact			Impact needed to eliminate waitlist over time frame
		5%	15%	25%	
Heart	11%	8%	3%	-3%	207%
Kidney	83%	76%	59%	42%	106%
Liver	-4%	-9%	-17%	-26%	109%
Lung	-56%	-58%	-63%	-67%	97%
Pancreas	-21%	-22%	-24%	-26%	373%
Combined	54%	48%	34%	20%	373%

4.2.2.3 Survival benefit

Using the estimated survival benefit associated with organ transplant, Table 4.4 shows the expected life years gained under varying impact levels with a random and ideal allocation. Our model shows that over the study period, the presumed consent base case would have gained an annual average of approximately 4,000 life-years under the random allocation scenario, with the biggest impacts seen in kidney and liver transplants. With a 25% impact, this increased up to 21,000 life-years gained. Under the ideal allocation scenario, the average annual life-years gained ranged from 11,000 in the base case to over 34,000 at higher presumed consent policy impacts.

4.2.3 Discussion

Presumed consent is a policy that has been proposed in several countries worldwide, however the realistic impact of such a policy in the US are uncertain. In our analysis we used estimates of the impact of presumed consent derived from other countries that have instituted such a policy worldwide. [77, 82] We found that even in the best case scenario, the presumed consent policy would have modest impacts on waitlists for all organs and would not resolve removals due to death or illness under an ideal allocation policy. This highlights the multifaceted approach that is needed even if deceased donation is increased through policy change such as presumed consent.

4.2.3.1 Expectations and precedents for presumed consent

In our analysis, we modeled a 5%-25% increase in deceased donation and conservatively used 5% as a base case because of the well-established deceased donor transplant system with relatively high donation rates in the US. Our upper bound of a 25% increase was based on the largest systematic review of presumed consent which included Austria, Belgium, and Singapore, where presumed consent resulted in a 20-25% increase in the number of donors available. [77] Data from before

Table 4.4: Average annual estimated life years gained by organ and presumed consent impact for both random and ideal allocations

Random Allocation				
Organ	Life Years per Transplant	Presumed Consent Impact		
		5%	15%	25%
Heart	4.4	3,452	10,357	17,262
Kidney	4.3	24	72	119
Liver	4.9	892	2,676	4,461
Lung	2.6	139	416	694
Pancreas	2.4	69	208	346
Combined	–	4,283	12,849	21,415

Ideal Allocation				
Organ	Life Years per Transplant	Presumed Consent Impact		
		5%	15%	25%
Heart	12.1	9,494	18,989	28,483
Kidney	14.5	144	288	432
Liver	10.1	2,095	4,191	6,286
Lung	9.5	269	538	807
Pancreas	4.9	130	261	391
Combined	–	11,355	22,710	34,066

and after implementation in several countries in South America and Europe have shown even more dramatic increases in liver and kidney donation rates (28%-1186%), however there are likely other policy changes and secular factors that may have influenced the results. [82] In contrast, recent data from the early experience in Wales, has shown no measurable impact on donor availability 18 months after the policy change. [67]

Several factors influence organ donation. Economic conditions, structural changes to the organ transplantation infrastructure, healthcare system characteristics, and societal norms must be taken into account when measuring the impact of policy changes over time. The mixed results following implementation of presumed consent policies likely reflect the complex interaction of these factors and highlight the difficulty of extrapolating the experience of one country to another. The specifics of a given country’s policy implementation may also impact the availability of deceased donors including the details on how patients opt-out and the ultimate role of the potential donor’s family in making organ donation decisions.

4.2.3.2 Public policy impact

One point of contention about implementation of presumed consent has been public backlash to such a policy. A 2012 survey of 3,200 representative US adults addressed presumed consent. The

survey found that the majority (51.1%) of surveyed American supported or strongly supported a presumed consent policy and 23.4% of patients would opt-out of such a system. [36] Additionally, early data from the policy implementation in Wales suggests that there has not yet been a public backlash to implementation. [67] However, careful planning and public education would be required prior to implementation of such a policy to ensure that the risk of backlash is minimized.

The other major potential drawback to a presumed consent policy could be a drop in living donor liver transplantation, currently a major source of kidney transplants in the US and a growing source for liver transplantation. However, as our analysis shows, a presumed consent alone would fail to counteract the ongoing growth of the waitlist for living donor eligible patients (kidney and liver) and thus there will be a continued need to increase donor availability in other ways, including promoting living donation.

4.2.3.3 Contextual factors

Our study has many strengths and limitations that warrant highlighting. First, we used a historical cohort to model the impact of a presumed consent system and these results may not necessarily apply to a contemporary cohort of patients. However, the population of patients with organ failure requiring transplant only continues to grow with commensurate growth in waitlist removal, thus the impact of our analysis may be conservative. [48] Other metrics such as organ yield may change with advent of new technologies (i.e. machine perfusion), and thus the overall impact of a policy on transplants and life years gained would likely fluctuate. Second, any organ allocation policy changes created for individual organs may impact the estimates presented in our analysis. Third, we could not accurately model the allocation of newly available organs associated with presumed consent, due to the central role in human decision making in organ allocation, thus we estimated a random and an ideal allocation system to approximate the lower and upper bounds. Finally, we derived the most appropriate estimate of life years gained currently available in the literature. However, in the random allocation policy, these estimates are a lower bound for life years gained, as they are based on already realized gains from transplant patients. Many of the patients in the study were still living at the time of the analysis where we derived the estimated benefits and were therefore still accruing life years. [76] Thus, these estimates for life years gained are also likely lower bounds for presumed consent impact. In the ideal allocation we used average post-transplant survival since all organs were allocated to patients who would have otherwise been removed from the waitlist.

4.3 Increasing Donation via “Ineligible” Donors

Variation in organ availability across OPOs is in part due to difference in organ acceptance patterns, including variation in the use of donors that do not meet OPTN eligibility criteria. [19, 27, 21] Eligible donors include people living within each OPO’s service area who died 75 years or younger with a body mass index (BMI) less than 50 kg/m², in addition to other organ-specific criteria. In general, donor eligibility criteria are designed to outline desirable health characteristics for organ donors. A formal eligibility definition is also helpful in understanding OPO and transplant center performance. For example, a key performance metric is the proportion of potential donors recovered and used for transplant within an OPO service area. However, the health characteristics of potential donor populations can vary among OPOs, and certain areas might have a particularly high prevalence of conditions that are prohibitive to the donation of certain organs. By measuring recovery metrics in terms of the proportion of eligible donors, we avoid penalizing centers for population characteristics out of their control. However, a drawback of the eligibility definition is that, in many cases, it only outlines desirable characteristics, not required characteristics. For example, there is a general age cut off of 75 years for a patient to be considered eligible, yet there are many examples of patients older than this cutoff who have successfully donated. Across OPOs, the use of ineligible donors varies significantly, with anywhere from 5% to 39% of deceased organ donations coming from ineligible donors. [21] Previous work has also found that donor eligibility plays a major role in our understanding of OPO performance and the state of the US transplant system more generally. [21] In 2020, the Advancing American Kidney Health executive order called for improved organ donation metrics, and the Centers for Medicare and Medicaid Services issued a final rule standardizing the eligibility definition based on potential donors’ cause of death. [68] A national policy to standardize the use of ineligible donors could spur efforts within low-performing OPOs to improve organ procurement and increase overall organ availability.

A barrier to increasing the use of ineligible donors are potential concerns about worsening transplant recipient outcomes (e.g. survival) when using “lower quality” organs. [64] Historically, survival outcomes after receiving an ineligible organ donation vary by organ, center, and patient subgroup. [64, 50, 61, 79, 15, 14, 46] This fact, combined with the significant differences in ineligible donor use by OPO, obscures the potential outcomes of a national policy to standardize ineligible donor use.

In this section, we perform a comprehensive analysis on the association of donor eligibility with graft and patient survival. Our goal was to offer insight on the viability and impact of increasing ineligible donor use. We combine national donation and survival data on solid-organ transplants to compare the outcomes of eligible and ineligible donations and to isolate this association. We also aimed to model these results as potential life-year increases associated with best practice sharing

across OPOs.

4.3.1 Modeling approach

Our analysis included data from the Standard Transplant and Research (STAR) files collected by the United Network of Organ Sharing (UNOS) from January 2008 through November 2020. UNOS is the private, non-profit organization that manages organ transplant in the United States under contract with the federal government. We included adult recipients of deceased donor solid organ transplants: heart, kidney, liver, lung, and pancreas. Our data also included UNOS-reported information categorizing eligible vs non-eligible deaths over the study period. Eligible deaths were patients declared brain dead according to state and local law with no exclusionary criteria as defined by OPTN policy. Example criteria include that prospective donors must be 75 years or younger with a BMI less than 50 kg/m^2 , in addition to other organ-specific criteria. A complete list of the eligibility criteria is included in Appendix C. For OPO-level analyses, each organization was given a unique identifier distinct from the STAR dataset. The University of Michigan institutional review board provided an exemption for the secondary use of de-identified data in this study.

4.3.1.1 Donor eligibility

The broad spectrum of eligibility characteristics for individual donors are not stored in the UNOS data set in their entirety. Instead, the overall eligibility of an individual donor is reported to UNOS. As a result, we defined an eligible donor as any donor listed in both the STAR file deceased donor dataset and the UNOS reported eligible death dataset. Likewise, an ineligible donor was any donor listed in the deceased donor dataset but not in the eligible death dataset. We compared the mean age and BMI of eligible and ineligible donors using t-tests. We calculated the percent of donors meeting eligibility requirements by organ, sex, ethnicity, and year of donation. We tested for significant differences across eligible and ineligible groups using chi-squared tests. Data preparation and statistical tests comparing populations were performed using Python version 3.7.9. Statistical significance was defined as a $p\text{-value} < 0.05$.

4.3.1.2 Association of donor eligibility and survival

We analyzed graft and patient survival by organ for both eligible and ineligible donors using Kaplan Meier curves. We compared survival rates using log rank tests. To better understand the influence of brain death versus cardiac death on eligibility-related survival, we compared Kaplan Meier curves for 1) the entire study population and 2) exclusively donation after brain death (DBD) donations as donation after cardiac death (DCD) donor status is a common reason for ineligibility.

We calculated hazard ratios for age, ethnicity, gender, BMI, blood type compatibility, donor cause of death (COD), and OPO of donation using a Cox proportional hazards model. We included interaction effects between donor eligibility and all main effects for kidney, liver, lung, and pancreas transplants. We did not include these interaction terms for heart transplants, as the relatively small number of ineligible heart donations prevented convergence for several regression parameters, making the model uninterpretable. Our regression tests were performed in R version 4.0.5.

4.3.1.3 Estimated impact of increasing ineligible donor use

We estimated the effect that increasing the use of ineligible donors would have had on the number of transplants and the number of life years gained over the study period. For each organ (heart, kidney, liver, lung, pancreas), we began by calculating the ineligible donor use rate by OPO. In those OPOs which had an ineligible donor utilization rate below pre-specified percentiles (50th, 75th, and 100th percentiles, respectively), we modeled scenarios where the OPOs would increase utilization to the pre-specified percentile. In the simulations, we did not change the ineligible donor use rate of OPOs with rates at or above the specified percentile. After simulating the adjusted ineligible donor use rates, we calculated the number of additional transplants corresponding to this increase in donations.

After calculating the increase in the number of transplants, we converted this increase into a life-years gained metric using data from Rana et. al (2015). [76] Rana et. al used propensity score matching to estimate the increase in life-years from receiving a transplant versus remaining on the waitlist. We leveraged this propensity score matching to improve our estimates of life-years gained. We adjusted Rana et. al's estimated life-years gained per transplant by the survival difference between eligible and ineligible donations found in our own model. We mapped the estimated increases in life-years gained using Tableau version 2020.4.

4.3.2 Numerical results

4.3.2.1 Overall cohort characteristics

From January 2008 through November 2020, there were 296,095 adult solid organ transplants (61% male, 54% White individuals). Table 4.5 provides a description of the study population. Over the study period, 86% of donors met eligibility requirements. Eligibility rates varied by organ, with as many as 20% of kidney donations and as few as 2% of heart donations coming from ineligible donors. Figure 4.6 shows the distribution of ineligible donor use rates across OPOs.

The distribution of eligible and ineligible donor groups was statistically different in terms of both donor and recipient sex and ethnicity ($p < 0.01$ for all). White individuals were more likely to

Table 4.5: Organ donor population description. P-values are the result of χ^2 -test (categorical) and t-test (continuous) variable comparisons between eligible and ineligible subgroups. SD – standard deviation.

	All donations	Eligible donations	Ineligible donations	p-value
N	297,223	255,039	42,184	
N by organ (%)				<0.001
Heart	30,103 (10%)	29,540 (12%)	563 (1%)	
Kidney	152,216 (51%)	121,240 (48%)	30,976 (73%)	
Liver	79,180 (27%)	70,102 (27%)	9,078 (22%)	
Lung	24,802 (8%)	23,598 (09%)	1,204 (3%)	
Pancreas	10,922 (4%)	10,559 (04%)	363 (1%)	
N by donor ethnicity (%)				<0.001
White	197,068 (66%)	163,965 (64%)	32,676 (78%)	
Black	45,979 (15%)	42,049 (16%)	3,860 (9%)	
Hispanic	41,948 (14%)	38,086 (15%)	3,804 (9%)	
Other	12,228 (4%)	10,939 (4%)	1,289 (3%)	
N by recipient ethnicity (%)				<0.001
White	164,338 (56%)	142,613 (57%)	21,725 (52%)	
Black	66,992 (23%)	56,408 (22%)	10,584 (25%)	
Hispanic	42,439 (14%)	36,195 (14%)	6,244 (15%)	
Other	19,729 (7%)	16,240 (6%)	3,489 (8%)	
N by donor sex (%)				<0.001
Male	183,766 (62%)	156,759 (61%)	27,007 (64%)	
Female	113,457 (38%)	98,280 (39%)	15,177 (36%)	
N by recipient sex (%)				0.0395
Male	187,645 (64%)	160,954 (64%)	15,351 (37%)	
Female	105,853 (36%)	90,502 (36%)	26,691 (63%)	
Mean donor age (yr) (SD)	39.38 (14.15)	38.82 (13.92)	42.75 (14.99)	<0.001
Mean recipient age (yr) (SD)	53.56 (12.53)	53.31 (12.58)	55.06 (12.09)	<0.001
Donor BMI (kg/m²)	27.84 (6.43)	27.68 (6.26)	28.76 (7.30)	<0.001
Recipient BMI (kg/m²)	27.96 (5.47)	27.86 (5.46)	28.56 (5.46)	<0.001

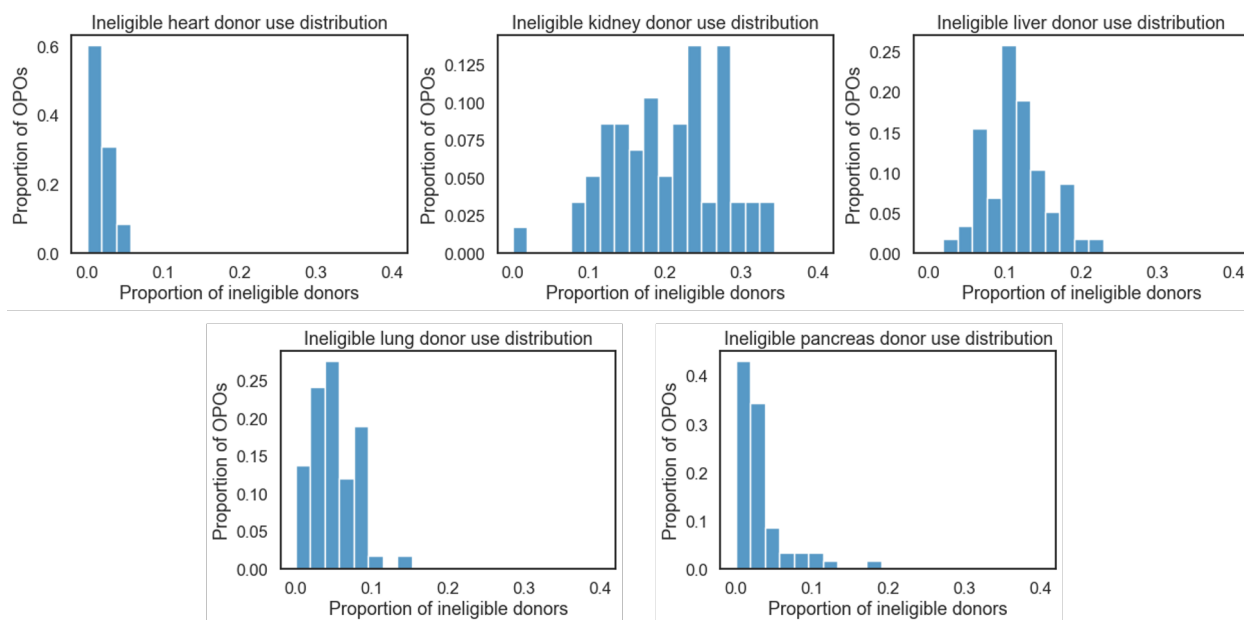


Figure 4.6: Histograms of ineligible donor use rates across OPOs, stratified by organ.

be an ineligible donor and less likely to receive a transplant from an ineligible donor, when compared to non-white individuals. Similarly, men were more likely than women to be an ineligible donor but less likely to receive a transplant from an ineligible donor.

4.3.2.2 Donor eligibility and survival

Figure 4.7 shows Kaplan Meier curves of graft survival for recipients of both eligible and ineligible donors by organ. There were statistically significant differences in graft survival between kidney and liver donations using eligible and ineligible donations ($p < 0.01$ for both). Recipients of ineligible kidney donations saw a 0.74%, 1.12%, 1.28%, and 2.20% relative decrease in 1-, 3-, 5-, and 10-year graft survival probability, respectively. Recipients of ineligible liver donations saw a 3.62%, 4.99%, 5.88%, and 9.38% relative decrease in 1-, 3-, 5-, and 10-year graft survival probability, respectively.

Figure 4.8 shows Kaplan Meier curves for patient survival. Log rank tests also found statistically significant differences in patient survival for kidney and liver donors ($p = 0.01$ and $p < 0.01$, respectively). Recipients of ineligible kidney donations saw a 0.21%, 0.43%, 0.45%, and 1.48% relative decrease in 1-, 3-, 5-, and 10-year patient survival probability, respectively. Recipients in ineligible liver donations saw a 1.35%, 2.37%, 3.15%, and 7.22% decrease in 1-, 3-, 5-, and 10-year patient survival probability, respectively. There were no statistically significant differences in patient survival for heart, lung, and pancreas transplants using ineligible donors, when compared to transplants using eligible donors.

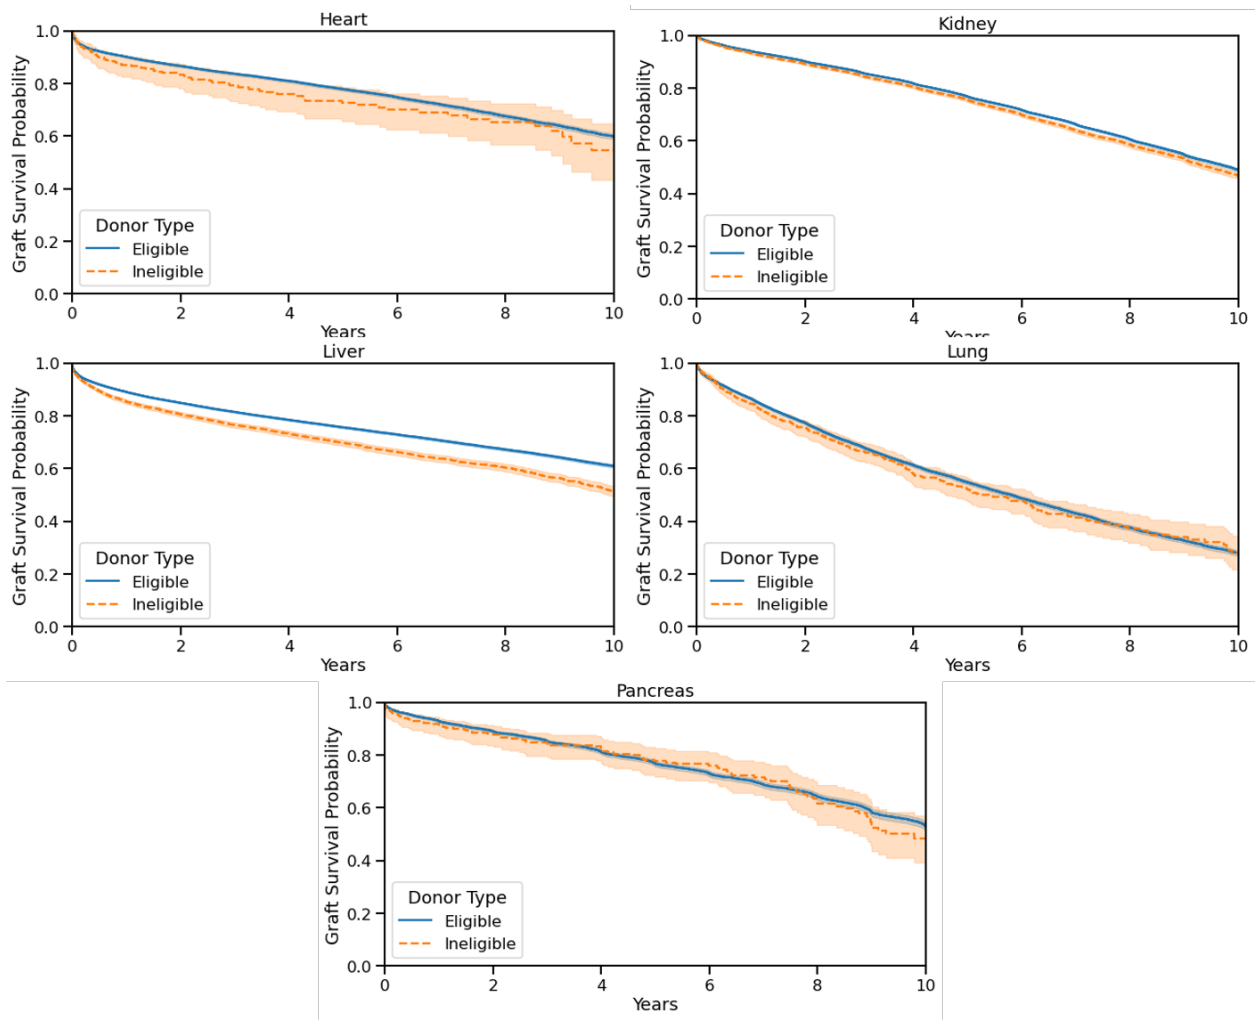


Figure 4.7: Kaplan Meier curves for transplant graft survival from January 2008 through November 2020, stratified by organ and donor eligibility.

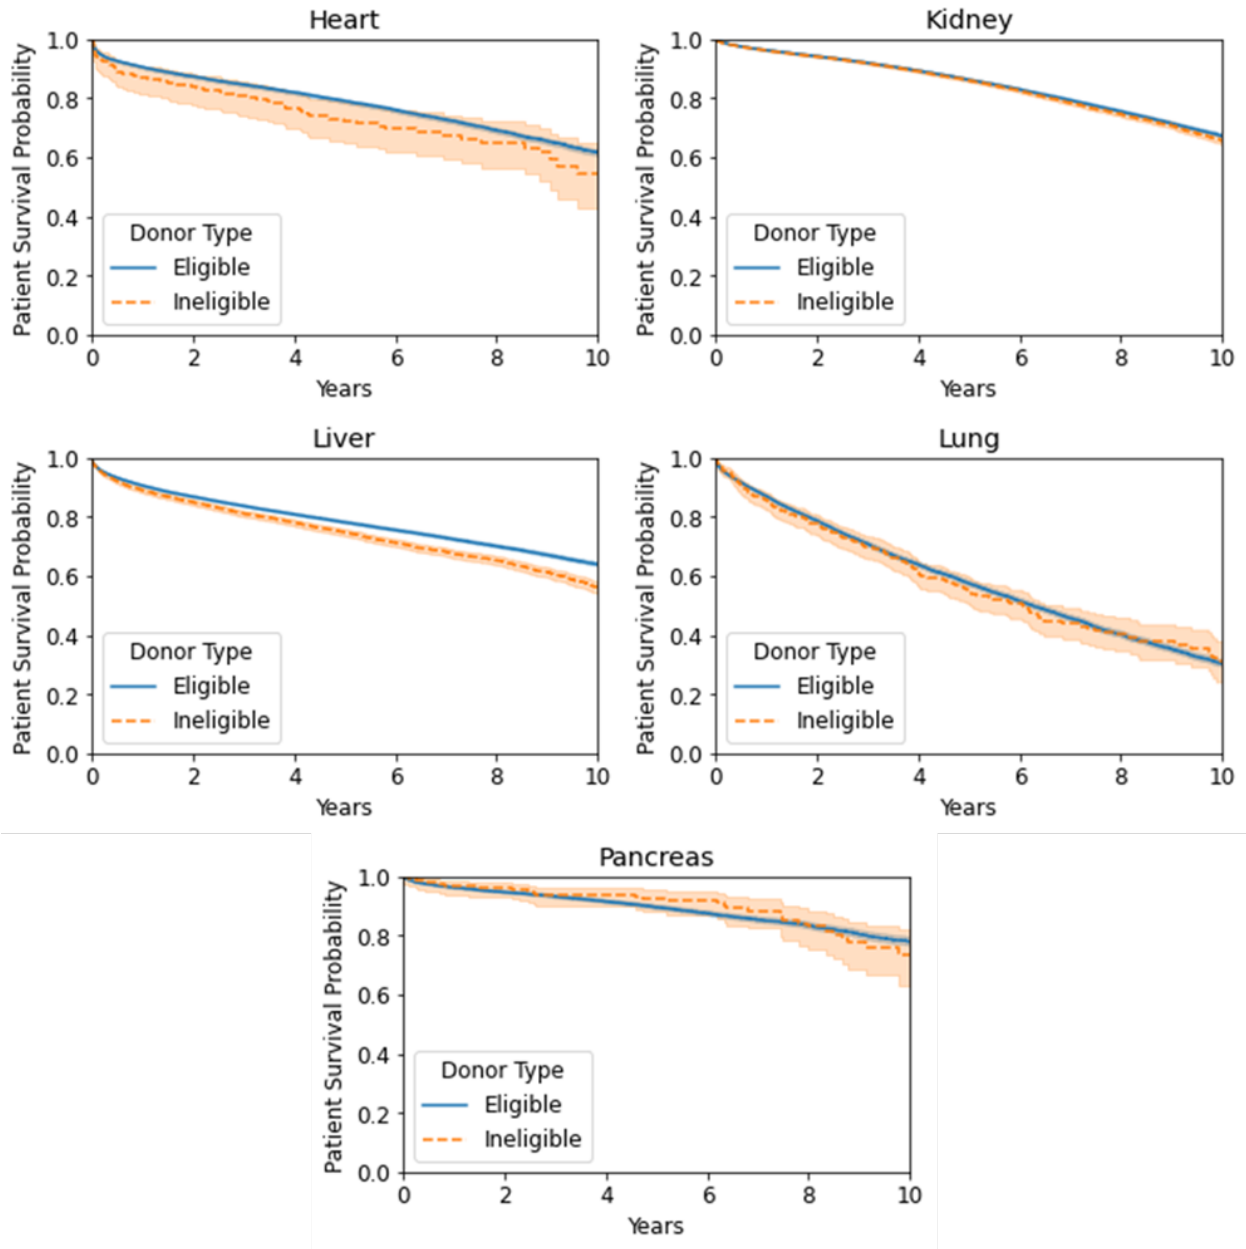


Figure 4.8: Kaplan Meier curves for transplant patient survival from January 2008 through November 2020, stratified by organ and donor eligibility.

4.3.2.3 Donation after cardiac versus brain death donor eligibility and survival

Figures 4.9-4.11 show Kaplan Meier curves by DBD and DCD status. Log rank tests found that, across the entire population, DCD donations provided lower graft survival for kidney and liver transplants ($p < 0.01$ for both) and no statistically significant difference for other organs. However, when exclusively examining ineligible donations, ineligible DBD donations provided significantly lower graft survival for kidney and pancreas donations compared to eligible DBD donations ($p < 0.01$ and $p = 0.01$, respectively) with no statistically significant differences for other organs. Our analysis also found that survival loss was primarily associated with ineligible DBD donors, and not DCD donors. When looking exclusively at DBD donations across all organs, the 10-year graft survival probability loss associated with ineligible donors increased by 6.90% when compared to all donors.

4.3.2.4 Multivariate survival modeling

A Cox proportional hazard model for each organ type provided estimates of the association of graft survival with donor eligibility and OPO, as well as recipient age, sex, ethnicity, and BMI. Additionally, the models provided estimates for the interaction effects of each variable with donor eligibility. Figure 4.12 shows hazard ratio (HR) estimates for the main effect of each variable by organ, excluding the OPO variable. Due to the large number of OPOs, hazard ratios for the OPO of donation are shown by organ in Figure C.1 in Appendix C. After accounting for demographic factors, likelihood ratio tests found that donor eligibility (including its interaction effects with other variables) was significantly associated with kidney, liver, and lung graft survival (kidney, $p = 0.01$; liver, $p < 0.01$; lung, $p = 0.02$;). Figure 4.13 shows HR estimates for the interaction of eligibility and all other variables.

4.3.2.5 Eligibility interactions

Recipient age and ethnicity were significantly associated with graft survival for all organs (Age: heart, $p < 0.01$; kidney, $p < 0.01$; liver, $p < 0.01$; lung, $p < 0.01$; pancreas, $p = 0.04$. Ethnicity: heart, $p = 0.04$; kidney, $p < 0.01$; liver, $p < 0.01$; lung, $p = 0.03$; pancreas, $p < 0.01$). Interaction terms suggested that the survival loss associated with ineligibility was significantly increased for older recipients of kidney and lung transplants ($p < 0.01$ for both). Two significant interactions between eligibility and ethnicity were found: one in Black recipients of kidney transplants, with a reduced HR effect of 0.93 (95% CI: 0.87-0.9873, $p = 0.02$); and another in lung transplant recipients of Other ethnicities, with an increased HR effect of 2.4520 (95% CI: 1.30-4.62, $p < 0.01$).

Sex was significantly associated with graft survival for kidney, liver, lung, and pancreas transplants (kidney, $p < 0.01$; liver, $p < 0.01$; lung, $p < 0.01$; pancreas, $p = 0.02$), however there was no

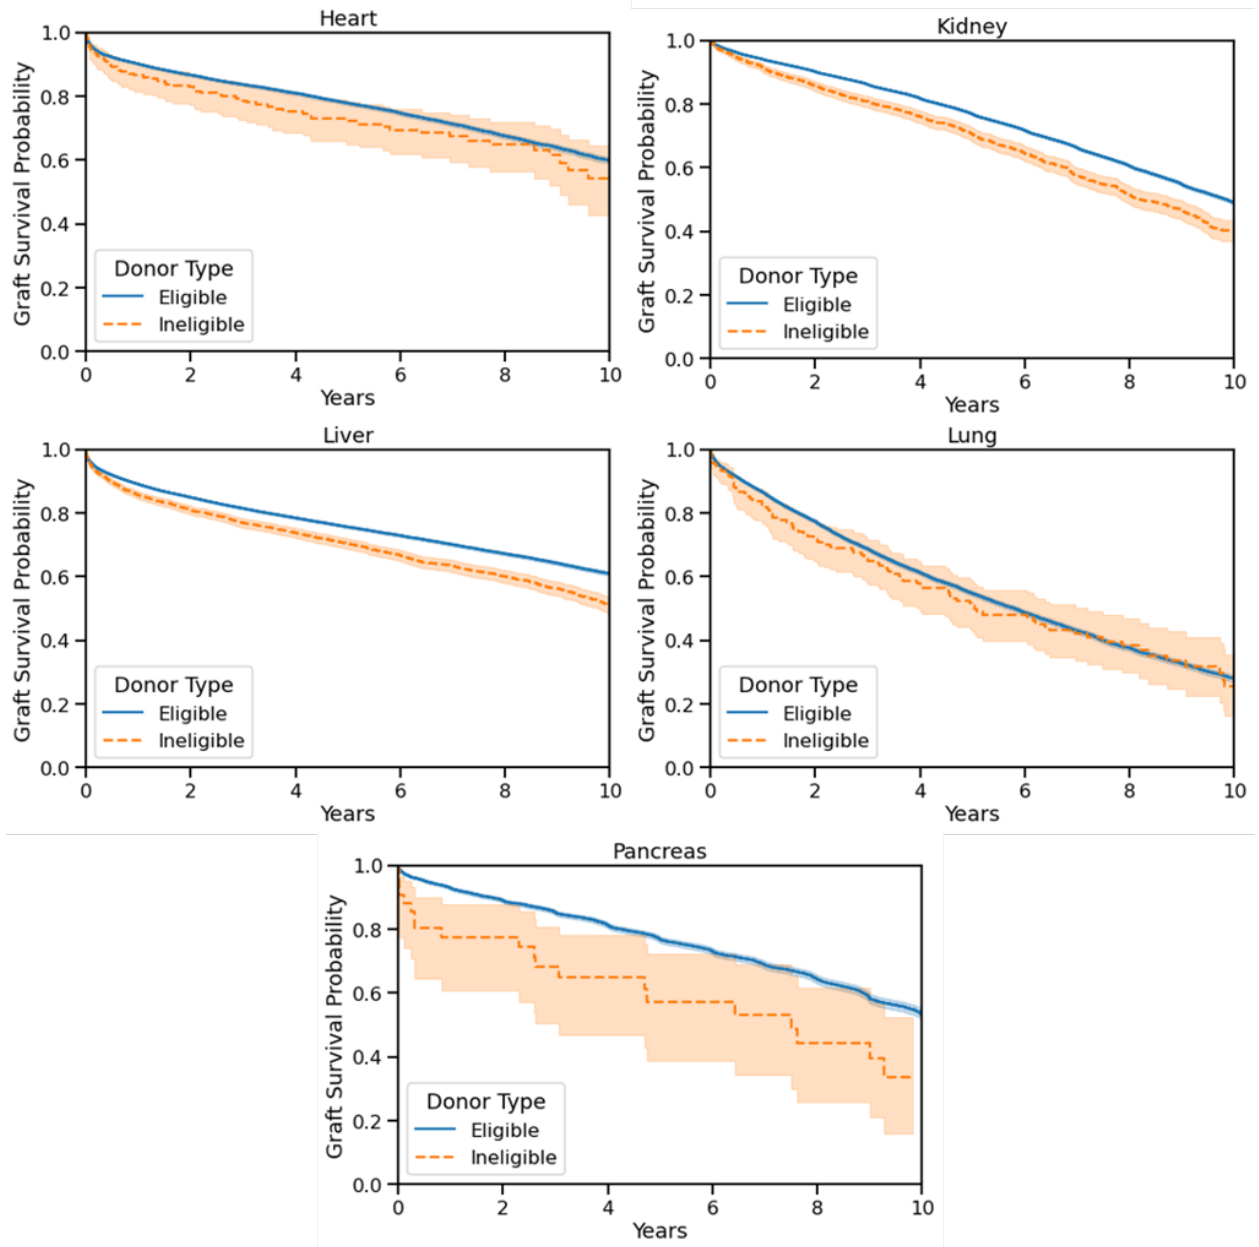


Figure 4.9: Kaplan Meier curves for post-transplant graft survival of deceased brain death donations, stratified by organ type and donor eligibility.

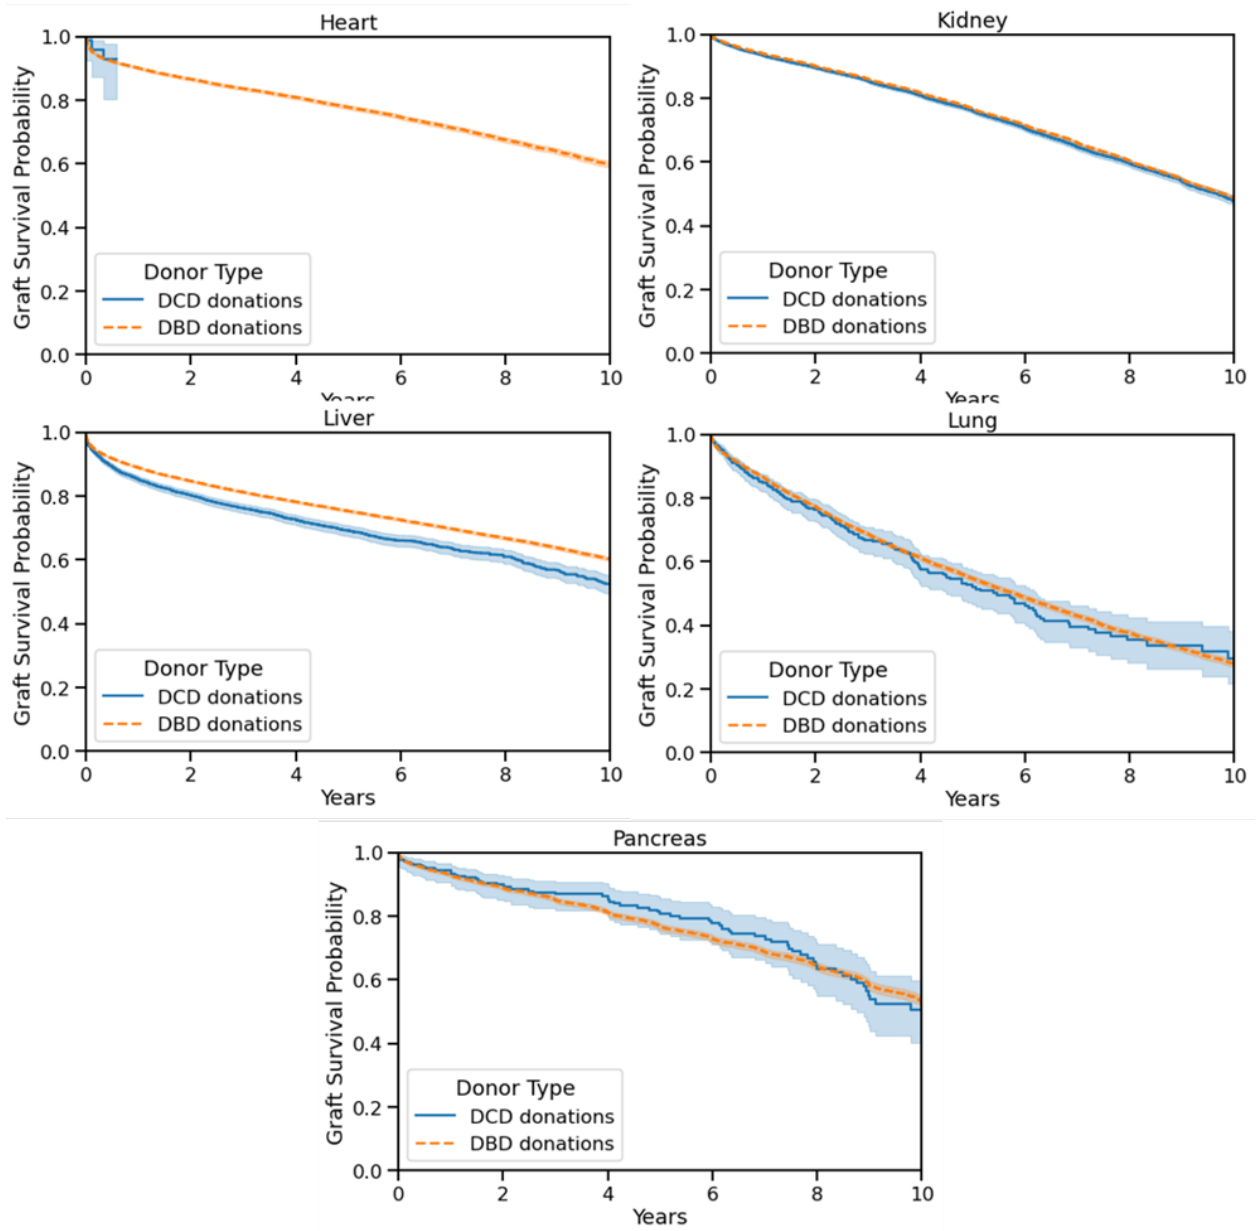


Figure 4.10: Kaplan Meier curves for post-transplant graft survival, stratified by organ type and donor death type (DCD: deceased cardiac death, DBD: deceased brain death).

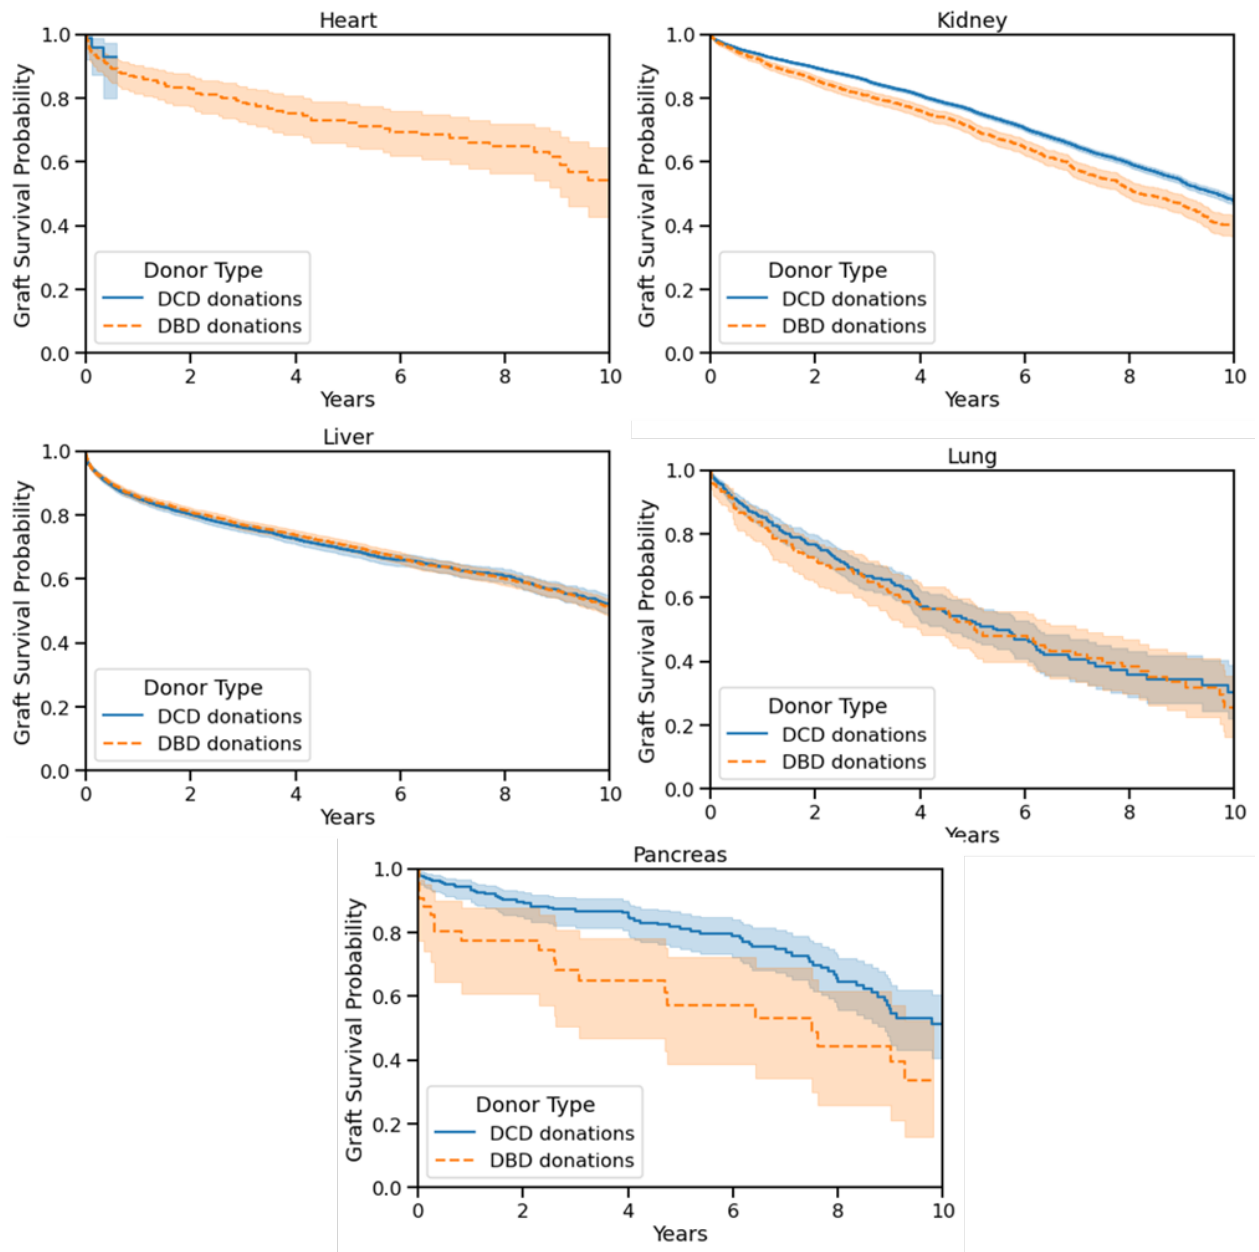


Figure 4.11: Kaplan Meier curves for post-transplant graft survival of ineligible donations, stratified by organ type and donor death type (DCD: deceased cardiac death, DBD: deceased brain death).

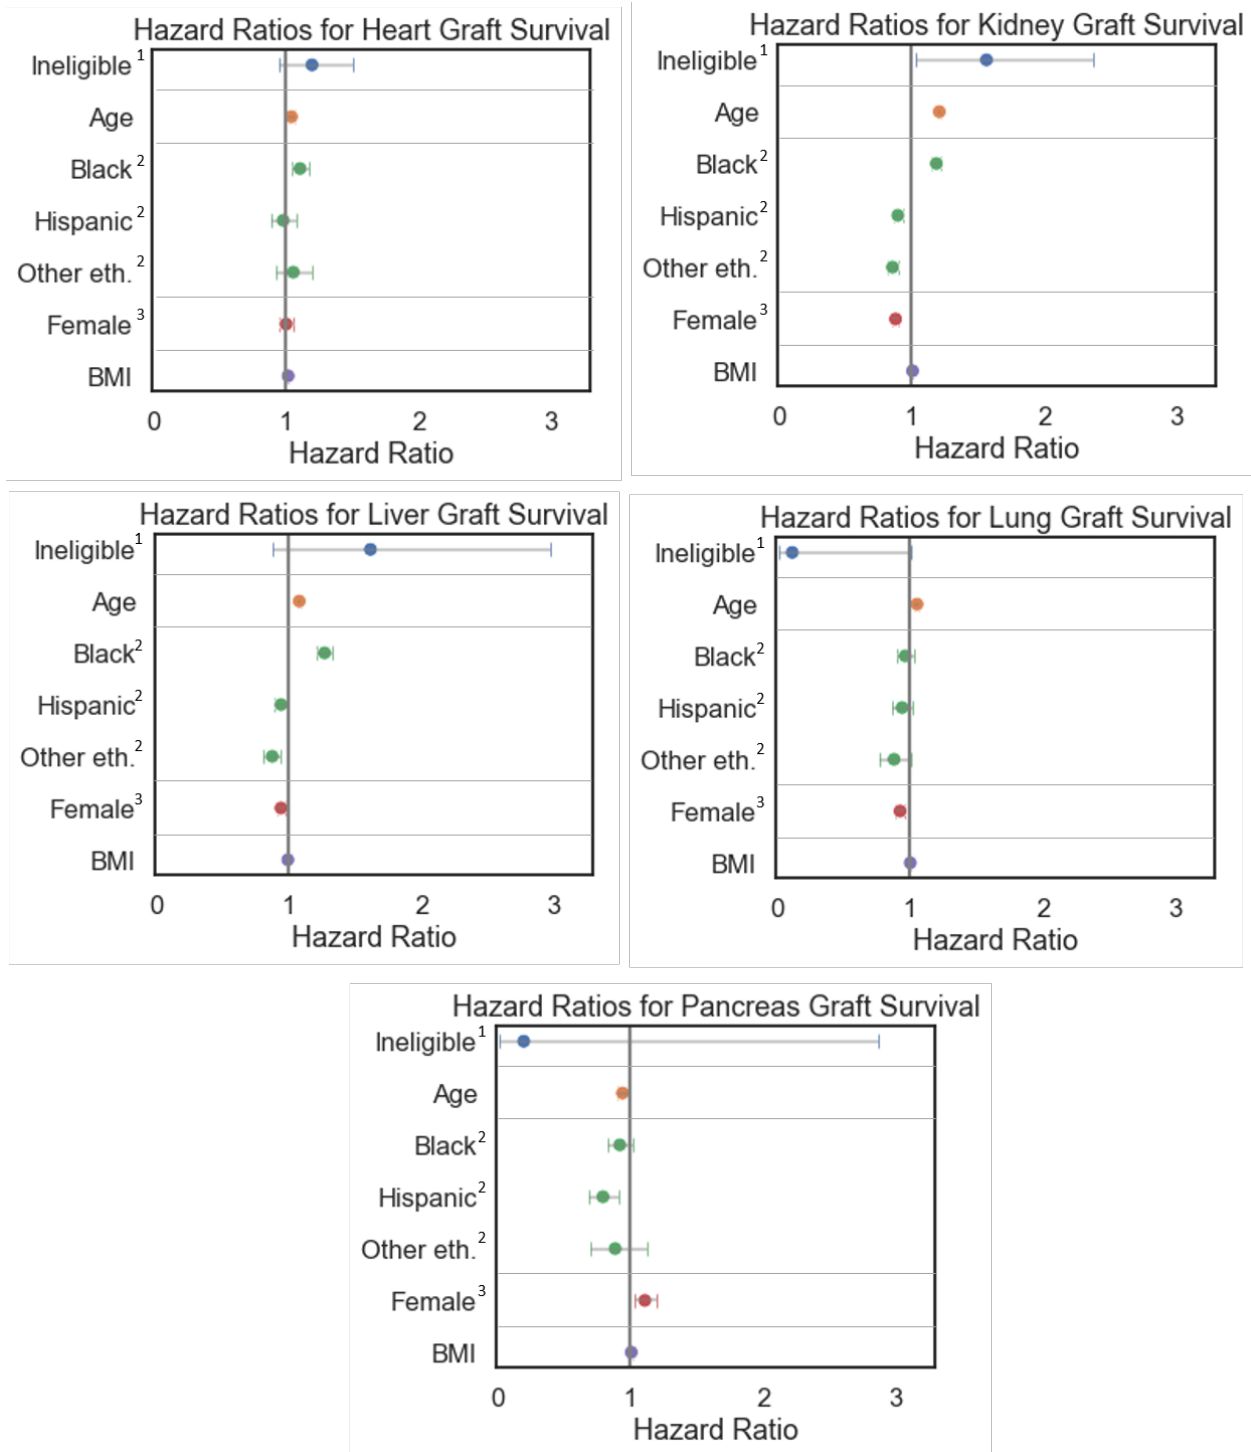


Figure 4.12: Selected hazard ratios calculated via Cox regression for transplant graft survival from January 2008 through November 2020. Recipient age is scaled to be in decades. Recipient age, ethnicity, and BMI were all significantly associated with graft survival ($p \leq 0.05$). Abbreviations: eth.: ethnicity.

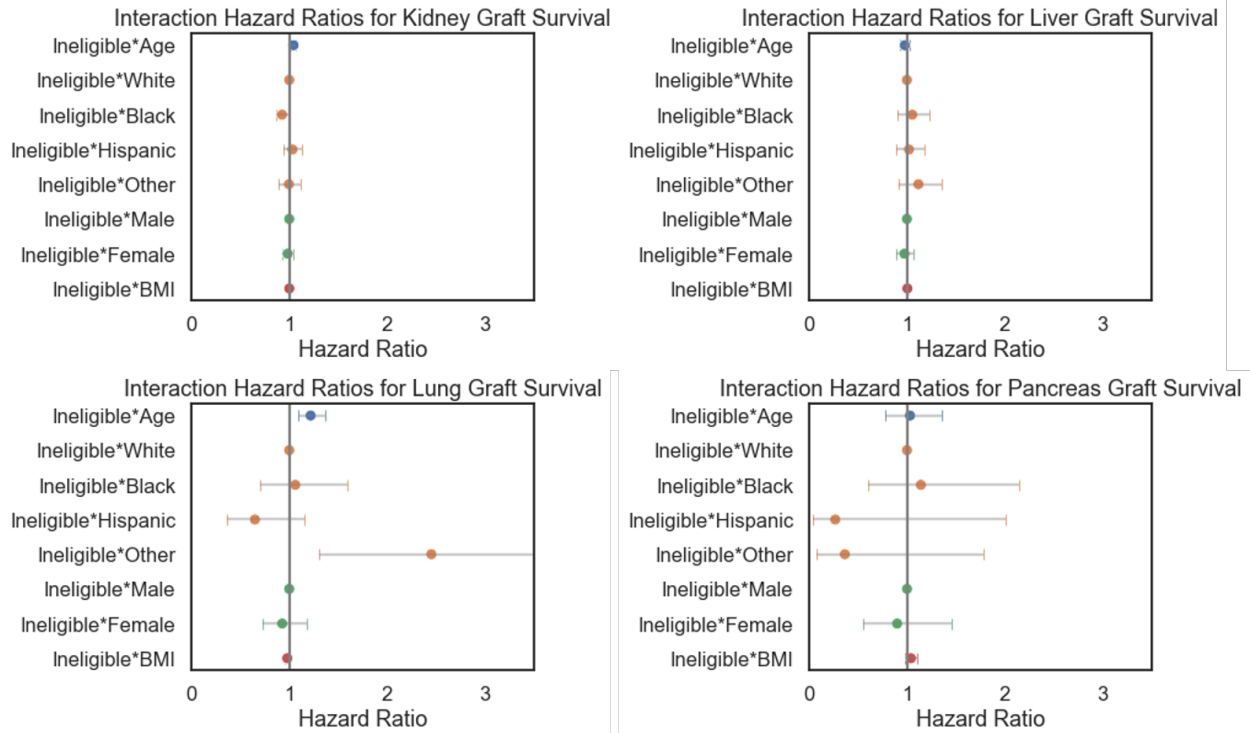


Figure 4.13: Selected hazard ratios calculated via Cox regression regarding the interaction effect of donor ineligibility with recipient age, ethnicity, sex, and BMI. Results are stratified by organ.

significant interaction found between donor eligibility and recipient sex. Similarly, while BMI was associated with graft survival for all organs (heart, $p < 0.01$; kidney, $p < 0.01$; liver, $p = 0.03$; lung, $p = 0.03$; pancreas, $p = 0.02$ for all), there was no statistically significant interaction found between donor eligibility and recipient BMI.

The OPO of donation was associated with graft survival for all organs ($p < 0.01$ for all). There were statistically significant interactions with OPO and eligibility which are not listed here for brevity.

4.3.2.6 Estimated impact of increasing ineligible donor use

We simulated the impact of individual OPOs increasing ineligible donor use rates to match the 50th, 75th, and 100th percentiles of use across the study period. For example, across all OPOs, the 75th percentile of ineligible liver donor use was 13.7%. Any OPO that had an ineligible donor use rate lower than 13.7% had its rate increased to this value in the 75th percentile match scenario. We did not change the ineligible donor use rate of any OPO having a rate greater than 13.7%. We then calculated the corresponding increase in donations and transplants with this increased donation rate. We performed this same technique for all organs and the other percentiles. Table 4.6 shows

the estimated mean increase in transplants and life-years gained for each percentile. Across all organs and OPOs, the estimated increase in the number transplants ranged from 6,061 to 33,470 throughout the study period, depending on the percentile. This translated to an estimated increase of 38,409 to 206,741 life-years gained over the same time frame.

Table 4.6: Estimated annual increases in transplants and life-years gained associated with increasing ineligible donor use under a range of scenarios. To calculate the percentile match results, all OPOs with ineligible donor use rates below the given percentile had their rates increased to match the percentile.

	Estimated life-years gained per transplant	Annual increase in # of transplants			Annual increase in life-years gained		
		50th percentile match	75th percentile match	100th percentile match	50th percentile match	75th percentile match	100th percentile match
Heart	6.8	7	21	89	49	145	603
Kidney	6.6	377	752	1,623	2,457	4,907	10,588
Liver	6.5	67	176	586	404	1,058	3,516
Lung	2.8	15	39	172	41	109	480
Pancreas	6.7	3	11	122	22	73	819
Overall	–	469	1,000	2,591	2,974	6,291	16,006

Figure 4.14 maps the increase in life-years gained by OPO if all OPOs below the 75th percentile rose to meet that use rate. At this percentile match, the overall increase in life-years gained ranged from 0 to 8,284 across the study period, depending on current ineligible use rate and the volume of transplants. OPOs shown with no increase in life-years gained were either at or above the 75th percentile of ineligible donor use.

4.3.3 Discussion

Reducing regional heterogeneity in the use of ineligible organ donations is one strategy to address the current donation shortage. For heart, lung, and pancreas transplants, using an organ from an ineligible donor had no difference in either graft or patient survival. In contrast, kidney and liver transplants had statistically lower survival rates when using ineligible donors. The rate of ineligible donor use varies from 5-39% across OPOs, suggesting that measures to better standardize ineligible donor use could be taken nationally. [21] Our modeling showed that increasing use of ineligible donors could result in an additional 469-2,591 organ transplants annually, providing an additional 2,974-16,006 life-years for waitlisted patients each year.

In December 2020, the Centers for Medicare and Medicaid Services issued a final rule that would update the eligibility definition to be based on standardized data from the Centers for Disease Control and Prevention mortality files (i.e. ICD-10 cause of death codes consistent with organ donation). [68] This update occurred after our study period, but could be an opportunity

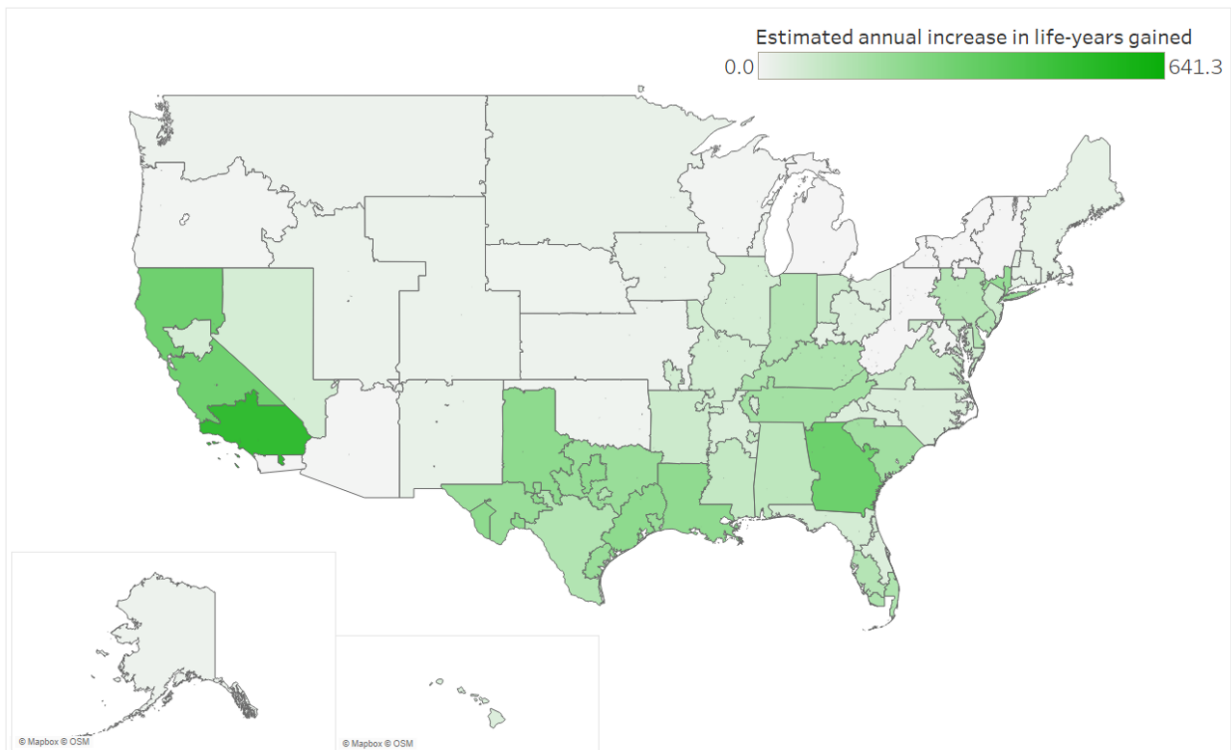


Figure 4.14: Map showing the estimated annual increase in life-years gained by organ procurement organization (OPO), if all OPOs with ineligible donation use rates below the 75th percentile increased their use to meet the 75th percentile. The increase shown is across all organs, based on annual use from January 2008 through November 2020.

to reduce regional heterogeneity in ineligible donor use. As this update is designed to reduce subjectivity in OPO performance metrics, it also provides an opportunity to evaluate differences in ineligible donor use across OPOs, with close attention and efforts to improve performance of under-performing OPOs. [30]

4.3.3.1 Survival differences for ineligible donations

The lack of significant detriment in graft survival for recipients of heart, lung, and pancreas ineligible donations may be related to low use rates (3-5%) of ineligible donations for these organs and stringent selection. Kidney and liver recipients, however, did see small but statistically significant decreases in graft survival when using ineligible donations. Notably, the median graft and patient survival for recipients of ineligible donations is still significantly higher than those of waitlisted patients who never receive a transplant, suggesting that there is a benefit to using select ineligible donations compared to no transplant at all. [76] We also note that a decrease in patient survival could have an impact on the timing of re-transplantation for recipients of ineligible donations. This suggests that perhaps the use of ineligible donations might be more appropriate in certain recipient populations. For example, older patients who may have a lower overall likelihood of requiring re-transplantation would also have a lower likelihood of being affected by any graft survival detriment derived from using an ineligible donation. However, the graft and patient survival loss from using ineligible donations was significantly higher for older recipients, so careful consideration would be required to balance these competing effects.

The data suggest that the graft survival detriment associated using ineligible kidney donations is primarily from the use of DBD donations. When looking exclusively at DBD donations across all organs, the 10-year graft survival probability loss associated with ineligible donors increased by 6.90% when compared to all donors. Additionally, while DCD donations were associated with worse graft survival for kidney and liver donations overall, this effect was actually reversed when considering only ineligible donations. Kidney and pancreas DCD donations actually corresponded with higher survival probabilities than ineligible DBD donations. Liver, heart, and lung DCD donations did not see a significantly lower graft survival probability than ineligible DBD donations. This could reflect a more stringent selection process of DCD donors.

4.3.3.2 Association of recipient demographics with survival

Our Cox regression model offers insight into the association of eligibility with graft survival when controlling for a recipient's age, ethnicity, sex, BMI, and OPO of donation. After controlling for these demographics, donor ineligibility and/or its interaction with recipient ethnicity was associated with lower graft survival for kidney, liver, and lung transplants.

Our Cox regression found higher hazard rates for Black recipients of heart, kidney, and liver donations, when compared to White recipients. Research suggests a number of potential sources for this disparity. Across organs, research suggests that a correlation between Black ethnicity and low socioeconomic status (SES) reduces access to medical care, and may result in difficulty adhering to treatment regimens. [56, 58, 93] However, prior studies have also shown that SES and adherence does not fully describe this disparity, and suggest that there are other systemic factors also contributing to this problem which are hard to define. [56, 58] Regarding kidney and liver transplants, Black recipients had higher a median time on dialysis prior to transplant, and a higher median model of end-stage liver disease (MELD) score at the time of listing, which are also correlated with decreased graft survival. [56, 93, 60]

Our model also found that the survival loss for kidney and liver transplants associated with the use of ineligible donors was increased for older recipients. However, recipients with higher BMI values did not experience an increased survival loss. This suggests that while older recipients have less reserve to handle marginal organs, high BMI patients might not have this same issue with reserve.

4.3.3.3 Potential Benefit of Increased Ineligible Donor Use

In many OPOs, the use of ineligible donations is already a common tool for addressing the need to increase organ donation. [21] However, the use of ineligible donations varies drastically by OPO. Our results suggest that the use of ineligible donations can provide a significant benefit to patients on the waiting list, and are a viable method for reducing patient morbidity and mortality. Regulatory changes in the definitions of eligibility, improved benchmarking of minimum standards of ineligible donor use, and/or sharing best practices to increase ineligible donor usage could improve ineligible donor utilization.

4.3.3.4 Contextual factors

This study benefitted from a large sample size via the use of a national, multi-organ data set containing nearly 13 years of data. Previous studies on the survival effects of ineligibility have focused on individual organs and/or a limited number of transplant centers. By including all adult solid-organ transplants from January 2008 through November 2020, we were able to provide a holistic view of the impact of ineligible donor use in the US. Our use of recipient demographics and Cox regression allowed us to better understand the specific association of ineligibility after controlling for known correlates with survival. The broad scope of our analysis prevented us from including more detailed patient health characteristics, which might better describe ineligible donor use patterns. Specific reasons for donor ineligibility by organ were not readily available for a more granular

analysis of which factors that determine ineligibility lead to worse outcomes. Additionally, the eligibility definition was changed in January 2017, which directly affected OPO performance metrics and could have impacted ineligible donor use. [21] While the data do not suggest a sudden shift in use associated with this specific change, we used an eligibility definition based on a donor's eligibility at the time of donation to be as consistent as possible in our analyses. Small sample sizes for ineligible heart, lung, and pancreas in certain OPOs also made parameter estimation challenging in our Cox regression. As a result, we were unable to estimate interaction effects between donor eligibility and some recipient characteristics for these organs.

We only considered graft and patient survival as quality outcomes when comparing eligible and ineligible donations. There are many other metrics that also provide valuable information on organ quality, such as delayed graft function, graft loss, readmission rates, etc., but were not within the scope of this analysis. Additionally, like most elements of the organ donation process, ineligible donor use depends on decisions made by individual clinicians, recipients, and other stakeholders, and is not fully dependent on OPOs. However, the granularity of our data limited our analysis to the OPO level. We also note that there may be heterogeneity in eligibility reporting standards across OPOs, which could cause misclassification in terms of eligibility and affect our calculation of donation rates and survival estimates by OPO.

4.4 Discussion

An immediate method to reduce the morbidity and mortality of patients awaiting an organ transplant is to increase the number of donated organs. In this chapter, we analyze the potential benefits of increasing organ donation via 1) implementing a presumed consent donation policy and 2) standardizing the use of ineligible organ donors.

In Section 4.2, we show the potential realistic impact of a presumed consent policy on solid organ transplant in the US. While there are ongoing ethical concerns regarding the implementation of such a policy that must first be resolved, a presumed consent policy is an immediate way to increase organ donation. We show that while presumed consent alone may not solve the organ shortage in the US it might result in large gains in life years and would be highly cost-effective for the US healthcare system. We believe that further public discourse about the ethics, logistics, and risks of a presumed consent policy are appropriate given the potential benefits.

In Section 4.3, we performed a national level review of graft and patient survival outcomes for ineligible versus eligible solid organ donations. We found that transplants using ineligible donations offer significant benefit for waitlisted patients who might otherwise never receive a transplant. We demonstrate that increasing the use of ineligible donors is an immediate method of reducing the current organ donation shortage and could significantly reduce patient morbidity and mortality.

We demonstrate that either policy has significant potential to reduce the gap between organ supply and demand and consequently, save the lives of thousands of patients with organ failure. We hope that this work provides insight to policy makers on the viability of these policies, and ultimately leads to improved outcomes for patients facing the immense burden of organ failure.

CHAPTER 5

Conclusion

Managing chronic conditions is one of the biggest healthcare challenges we face today. Over half of Americans have a chronic disease, and managing these conditions accounts for the vast majority of U.S. healthcare spend. [16] Because of the extreme burden of these conditions, the availability of resources like time, money, and medication play a large role in medical decision making. In this work, we present operations research methodologies to both reduce resource burden and increase resource availability.

In Chapter 2 we provide an MDP approach to identifying the optimal treatment interval for a chronic disease. By building a model that decides *when* to treat instead of *whether* to treat, we can derive closed-form solutions. These closed-form solutions allow for interpretable policy recommendations and can be presented to clinical decision makers as a menu of treatment options. As an application to AMD, we demonstrate that following this approach could save up to \$61 million in treatment costs annually.

In Chapter 3 we provide a dynamic programming approach to coordinating the treatment of multiple chronic conditions. Even if the optimal treatment interval for individual conditions is known, simultaneously considering multiple conditions adds significant complexity to treatment scheduling. In many cases, patients can significantly reduce their visit costs by coordinating multiple treatments at a single visit. However, this often comes at the expense of additional treatment costs. Beyond just providing a method for finding the optimal treatment schedule, we identify conditions under which it is always optimal to synchronize treatment. We also identify the specific synchronization policy under these conditions. For situations where these conditions are not met, we provide a simple, closed-form heuristic policy and bound the regret of following this policy. As an application to AMD, we demonstrate that following the optimal policy (and in fact, following the heuristic policy) can save \$582 million in direct medical costs alone.

In Chapter 4 we analyze and highlight two potential policies for increasing organ donation—a critical and limited resource for patients with chronic organ failure. For the 100,000+ patients on the organ waiting lists, we show that a presumed consent donation policy could add as many as

34,000 life-years annually. Similarly, we show that standardizing the use of ineligible organ donors could add up to 16,000 life-years annually.

5.1 Future Work

There are several avenues for future work building upon the ordinal MDP in Chapter 2. Examples of interesting extensions include:

- What is the optimal policy if the MSTI changes during the exploration phase?
- How does the exploration process change as the reward structure changes?
- How do patient and clinician risk tolerance affect decision making?

Some of these questions are relaxations of our ordinal MDP assumptions. One approach to relaxing the stationary MSTI assumption could assign a distribution to the outcome of each potential scheduling interval. Then, similar to our problem, should a clinician observe that an interval was too long (or short), she could update her distributional beliefs for that action as well as all actions longer than (or shorter than) the previously selected interval. While our model does allow for general reward structures, our numerical analysis and Theorem 2.4.2 are based on a constant per-period reward function. The behavior of the model in practice would likely vary as the reward structure takes on more complex forms, and is worth exploring. Additionally, we emphasize our model's ability to provide a menu of treatment options via its index policy. When provided with a list of options, some clinicians may decide to be more or less aggressive with treatment scheduling to either protect patient health or reduce patient cost. An interesting question involves how deviating from a risk-neutral optimization recommendation affects long-term patient outcomes.

The multiple condition model in Chapter 3 also has several interesting extensions. Some of the key questions include:

- What if the MSTIs for multiple conditions are unknown?
- What if clinicians are willing to schedule beyond the MSTI?
- How might multiple clinician decision makers affect the ability to synchronize treatment?

Perhaps the most immediate question in the multiple condition scenario involves what to do if the MSTIs are initially unknown. This extension would likely require the incorporation of health costs (such as exposure to fluid) in addition to direct medical costs, since we could no longer

feasibly guarantee that patients are not exposed to symptoms. For example, you could expand the model in Chapter 2 to consider multiple conditions. Considering the needs of multiple conditions during the exploration process could have a dramatic effect on the optimal searching policy, both in terms of search speed and patient health. For example, a patient may be willing to accept short-term symptoms from one of the conditions via a longer treatment interval if it provides useful information about the MSTIs of the other conditions.

A major practical consideration when managing multiple conditions is that patients are often treated by multiple healthcare providers. In these cases, synchronizing treatment is often not an easy logistical task. An interesting model could incorporate the concept of multiple decision makers. Potential avenues might include an additional synchronization cost, or a probability parameter around the likelihood that different treatments are actually able to be synchronized for a given visit. Additionally, it may be that conditions can be grouped according to which healthcare providers are managing them. For example, some combinations of conditions may be easily synchronized, while others are more difficult or even impossible.

The work presented in Chapter 4 can be most readily expanded by applying our population level findings to the patient level. The choice to use an organ donation from an ineligible donor is ultimately a decision between the individual patient and healthcare professionals. The choice is very complex and situational, and depends heavily on the individual patient, donor, hospital, and OPO. A decision making model for the individual patient could provide data-driven recommendations regarding organ acceptance, and a powerful next step in standardizing and improving ineligible donor use nationally. Some of the key questions that this model could address include:

- When is it preferable to accept an ineligible organ offer versus deferring for a potential future eligible organ offer?
- How do regional heterogeneity in donation rates and waitlist size affect individual acceptance decisions?
- How do acceptance decisions for one patient affect the remaining waitlist population?

As one example, this model could be formulated as an MDP, where the actions are to accept or reject an organ offer in a given period. The state space would likely include the patient's health over time, and the goal could be to maximize survival and quality of life. Even if such a model is unable to fully incorporate all of the context that goes into an acceptance decision, a data-driven framework could help identify and highlight which portions of the acceptance process can be improved.

5.2 Concluding Remarks

We hope that this work offers both theoretical and practical value to the field of medical decision making. Chronic diseases are both ubiquitous and incredibly personal. They effect the lives of virtually everyone, whether directly or through a loved one. By reducing the resource burden or increasing resource availability, we hope that the research presented here drives meaningful change and makes a lasting difference the lives of patients.

APPENDIX A

Supporting Material for Chapter 2

A.1 Algorithm for finding the optimal treatment interval

To help characterize the ordinal bandit, we provide a method for calculating the index values in Algorithm 1 below. The algorithm requires calculating and storing $\frac{|u-l| \cdot (|u-l|+1)}{2}$ cost-to-go values, one for each possible combination of lower and upper action indices.

Result: Index values $Q(d, P, l, u)$ and optimal decisions $a_{(P,l,u)}^*$

Initialize $P, l, u, r(\cdot|\cdot)$, and δ as appropriate;

for $d \leftarrow l$ **to** u **by** 1 **do**

 | Calculate $V(P, d, d) = \frac{r(d|d)}{1-\delta^{a_d}}$;

end

for $n \leftarrow 1$ **to** $u - l$ **by** 1 **do**

 | **for** $l \leftarrow 0$ **to** n **by** 1 **do**

 | **for** $d \leftarrow l$ **to** $l + n$ **by** 1 **do**

 | Calculate $Q(d, P, l, l + n)$;

 | **end**

 | $V(P, l, l + n) \leftarrow \max_{d \in \{l, \dots, l+n\}} \{Q(d, P, l, l + n)\}$;

 | $a_{(P,l,l+n)}^* \leftarrow \arg \max_{a_d \in \{a_l, \dots, a_{l+n}\}} \{Q(d, P, l, l + n)\}$;

 | **end**

end

APPENDIX B

Supporting Material for Chapter 3

B.1 Proof of Theorem 3.4.2

First, we define a cost-to-go value Φ^{π_1} that is the cost-to-go for following policy π_1 starting from a period in which both conditions must be treated (i.e. the cost-to-go when starting in state $X = (M_1, M_2)$). Under this policy, the clinician treats condition 1 every M_1 periods. Condition 2 is synchronized to be treated during the condition 1 visit that is closest to (but still less than) M_2 . Looking back at Figure 3.2, we note that it specifically shows a visual representation of the treatment cycle corresponding to Φ^{π_1} .

By following the cycle outlined in Figure 3.2 and noting that the visit and treatment costs represent geometric series, we can then write

$$\Phi^{\pi_1} = \frac{c_0 + c_1}{1 - \delta^{M_1}} + \frac{c_2}{1 - \delta^{M_1 \lfloor M_2/M_1 \rfloor}} \quad (\text{B.1})$$

Using Φ^{π_1} , we can also calculate the cost-to-go of following π_1 when starting in any state X . Since Φ^{π_1} is the cost-to-go when starting from a period in which both conditions are treated, calculating a general cost-to-go $V^{\pi_1}(x_1, x_2)$ requires us to first find the cost of getting to a period where we treat both conditions. This cost depends on how many periods are remaining before we need to treat each condition. We can consider 3 different scenarios:

- **Scenario 1:** The clinician must treat condition 1 multiple times before they will need to treat condition 2. Mathematically, this is when $(M_2 - x_2) - (M_1 - x_1) \geq M_1$.
- **Scenario 2:** The clinician must treat condition 1 exactly once before they will need to treat condition 2. Mathematically, this is when $(M_1 - x_1) \leq (M_2 - x_2) < M_1$.
- **Scenario 3:** The clinician must treat condition 2 before they will need to treat condition 1. Mathematically, this is when $(M_2 - x_2) < (M_1 - x_1)$.

In scenario 1, we must calculate the number of condition 1 visits required before the next condition 2 treatment. Under π_1 we can write this as:

$$\lambda(x_1, x_2) = \left\lfloor \frac{(M_2 - x_2) - (M_1 - x_1)}{M_1} \right\rfloor \quad (\text{B.2})$$

where the partial brackets represent the floor function.

In scenario 2, we will treat both conditions together at the next visit, which will occur when we need to treat condition 1 (i.e. in $M_1 - x_1$ periods). In scenario 3, we will again treat both conditions together at the next visit, which will instead occur when we need to treat condition 2 (i.e. in $M_2 - x_2$ periods).

Combining this information, we can write the cost-to-go under π_1 more generally as:

$$V^{\pi_1}(x_1, x_2) = \begin{cases} \sum_{i=1}^{\lambda(x_1, x_2)} \delta^{M_1 i - x_1} (c_0 + c_1) + \delta^{(\lambda(x_1, x_2) + 1)M_1 - x_1} \Phi^{\pi_1} & \text{if } (M_2 - x_2) - (M_1 - x_1) \geq M_1 \\ \delta^{M_1 - x_1} \Phi^{\pi_1} & \text{if } (M_1 - x_1) \leq (M_2 - x_2) < M_1 \\ \delta^{M_2 - x_2} \Phi^{\pi_1} & \text{if } (M_2 - x_2) < (M_1 - x_1) \end{cases} \quad (\text{B.3})$$

We next show the analytical results of policy evaluation across a range different state spaces. These ranges are based on the rules outlined in π_1 .

Case 1: $x_1 = M_1$ **and** $x_2 = M_2$

In this case, the clinician must treat both conditions in this period, as they are both at their MSTI values. As the only available action is to treat both conditions, π_1 is optimal regardless of c_0 . \square

Case 2: $x_1 < M_1$ **and** $x_2 = M_2$

In this case, the clinician must treat condition 2 in this period, but does not need to treat condition 1. The available actions are to treat both conditions or treat only condition 2. Under π_1 , the policy is to treat both conditions. We can use the following comparison to find necessary conditions for π_1 to be optimal.

Comparison 2.1: Treating both conditions must be at least as good as treating only condition 2.

Mathematically, this comparison is true when

$$c_0 + c_1 + c_2 + \delta V^{\pi_1}(1, 1) \leq c_0 + c_2 + \delta V^{\pi_1}(x_1 + 1, 1).$$

The truth of this comparison depends on the the current state X . We decompose this into two possible scenarios.

Subcase 2.1.1: $(M_2 - 1) - (M_1 - x_1 - 1) \geq M_1$

By substituting in for V^{π_1} and Φ^{π_1} and then isolating c_0 , we can rewrite this inequality as

$$c_0 \geq \frac{-c_1 * \left(\frac{\delta^{M_1-x_1} - 1}{1 - \delta^{M_1}} \right) - c_2 * \left(\frac{\delta^{(\lfloor \frac{M_2-M_1+x_1}{M_1} \rfloor + 1)M_1-x_1} - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}}{1 - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}} \right)}{\left(\frac{\delta^{M_1-x_1} - \delta^{M_1}}{1 - \delta^{M_1}} \right)} \quad (\text{B.4})$$

Subcase 2.1.1: $(M_2 - 1) - (M_1 - x_1 - 1) < M_1$

By substituting in for V^{π_1} and Φ^{π_1} and then isolating c_0 , we can rewrite this inequality as

$$c_0 \geq \frac{-c_1 * \left(\frac{\delta^{M_1-x_1} - 1}{1 - \delta^{M_1}} \right) - c_2 * \left(\frac{\delta^{M_1-x_1} - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}}{1 - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}} \right)}{\left(\frac{\delta^{M_1-x_1} - \delta^{M_1}}{1 - \delta^{M_1}} \right)} \quad (\text{B.5})$$

Therefore, we have found two equations which must be satisfied for π_1 to be an optimal policy. We will continue to do this for the remainder of the cases to see other relevant comparisons. \square

Case 3: $x_1 = M_1$ **and** $x_2 < M_2$ **and** $M_2 - x_2 < M_1$

In this case, the clinician must treat condition 1 in this period, but does not need to treat condition 2. The available actions are to treat both conditions or treat only condition 1. Under π_1 , the policy is to treat both conditions.

Comparison 3.1: Treating both conditions must be at least as good as treating only condition 1.

Mathematically, this comparison is true when

$$c_0 + c_1 + c_2 + \delta V^{\pi_1}(1, 1) \leq c_0 + c_1 + \delta V^{\pi_1}(1, x_2 + 1).$$

By substituting in for V^{π_1} and Φ^{π_1} and then isolating c_0 , we can rewrite this inequality as

$$c_0 \geq -c_1 - c_2 \left(\frac{\delta^{M_2-x_2} - 1}{1 - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}} \right) \left(\frac{1 - \delta^{M_1}}{\delta^{M_2-x_2} - \delta^{M_1}} \right) \quad (\text{B.6})$$

\square

Case 4: $x_1 = M_1$ **and** $x_2 < M_2$ **and** $M_2 - x_2 \geq M_1$

In this case, the clinician must treat condition 1 in this period, but does not need to treat condition 2. The available actions are to treat both conditions or treat only condition 1. Under π_1 , the policy is to treat only condition 1.

Comparison 4.1: Treating only condition 1 must be at least as good as treating both conditions.

Mathematically, this comparison is true when

$$c_0 + c_1 + \delta V^{\pi_1}(1, x_2 + 1) \leq c_0 + c_1 + c_2 + \delta V^{\pi_1}(1, 1).$$

By substituting in for V^{π_1} and Φ^{π_1} and then isolating c_0 , we see that this is true for all $c_0 \geq 0$.

□

Case 5: $x_1 < M_1$ **and** $x_2 < M_2$ **and** $M_1 - x_1 < M_2 - x_2$

In this case, the clinician does not have to treat either of the conditions in this period, and all actions are available. Condition 1 is closer to its MSTI than condition 2. Under π_1 , the policy is to treat neither condition.

Comparison 5.1: Treating neither condition must be at least as good as treating both conditions.

Mathematically, this comparison is true when

$$\delta V^{\pi_1}(x_1 + 1, x_2 + 1) \leq c_0 + c_1 + c_2 + \delta V^{\pi_1}(1, 1).$$

By substituting in for V^{π_1} and Φ^{π_1} and then isolating c_0 , we see that this is true for all $c_0 \geq 0$.

□

Comparison 5.2: Treating neither condition must be at least as good as treating only condition 1.

Mathematically, this comparison is true when

$$\delta V^{\pi_1}(x_1 + 1, x_2 + 1) \leq c_0 + c_1 + \delta V^{\pi_1}(1, x_2 + 1).$$

The above comparison depends on the current state X , so we separate this proof into 3 subcases.

Subcase 5.2.1: $(M_2 - x_2 - 1) - (M_1 - x_1 - 1) \geq M_1$

By substituting in for V^{π_1} and Φ^{π_1} and then isolating c_0 , we can rewrite this inequality as

$$c_0 \geq -c_1 - c_2 \left(\frac{\delta^{(\lfloor \frac{M_2 - x_2 - M_1}{M_1} \rfloor + 1)M_1} - \delta^{(\lfloor \frac{M_2 - x_2 - M_1 + x_1}{M_1} \rfloor + 1)M_1 - x_1}}{1 - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}} \right) \left(\frac{1 - \delta^{M_1}}{1 - \delta^{M_1 - x_1}} \right) \quad (\text{B.7})$$

Subcase 5.2.2: $(M_1 - 1) \leq (M_2 - x_2 - 1) - (M_1 - x_1 - 1) < M_1$

By substituting in for V^{π_1} and Φ^{π_1} and then isolating c_0 , we can rewrite this inequality as

$$c_0 \geq -c_1 - c_2 \left(\frac{\delta^{M_1} - \delta^{M_1 - x_1}}{1 - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}} \right) \left(\frac{1 - \delta^{M_1}}{1 - \delta^{M_1 - x_1}} \right) \quad (\text{B.8})$$

Subcase 5.2.2: $(M_2 - x_2 - 1) < (M_1 - 1)$

By substituting in for V^{π_1} and Φ^{π_1} and then isolating c_0 , we can rewrite this inequality as

$$c_0 \geq -c_1 - c_2 \left(\frac{\delta^{M_2-x_2} - \delta^{M_1-x_1}}{1 - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}} \right) \left(\frac{1 - \delta^{M_1}}{1 - \delta^{M_1} + \delta^{M_2-x_2} - \delta^{M_1-x_1}} \right) \quad (\text{B.9})$$

Comparison 5.3: Treating neither condition must be at least as good as treating only condition 2.

Mathematically, this comparison is true when

$$\delta V^{\pi_1}(x_1 + 1, x_2 + 1) \leq c_0 + c_2 + \delta V^{\pi_1}(x_1 + 1, 1).$$

By substituting in for V^{π_1} and Φ^{π_1} and then isolating c_0 , we see that this is true for all $c_0 \geq 0$.

□

Case 6: $x_1 < M_1$ **and** $x_2 < M_2$ **and** $M_1 - x_1 \geq M_2 - x_2$

In this case, the clinician does not have to treat either of the conditions in this period, and all actions are available. Condition 2 is as close or closer to its maximum safe treatment interval when compared to condition 1. Under π_1 , the policy is to treat neither condition.

Comparison 6.1: Treating neither condition must be at least as good as treating both conditions.

Mathematically, this comparison is true when

$$\delta V^{\pi_1}(x_1 + 1, x_2 + 1) \leq c_0 + c_1 + c_2 + \delta V^{\pi_1}(1, 1).$$

By substituting in for V^{π_1} and Φ^{π_1} and then isolating c_0 , we see that this is true for all $c_0 \geq 0$.

Comparison 6.2: Treating neither condition must be at least as good as treating only condition 1.

Mathematically, this comparison is true when

$$\delta V^{\pi_1}(x_1 + 1, x_2 + 1) \leq c_0 + c_1 + \delta V^{\pi_1}(1, x_2 + 1).$$

By substituting in for V^{π_1} and Φ^{π_1} and then isolating c_0 , we see that this is true for all $c_0 \geq 0$.

Comparison 6.3: Treating neither condition must be at least as good as treating only condition 2.

Mathematically, this comparison is true when

$$\delta V^{\pi_1}(x_1 + 1, x_2 + 1) \leq c_0 + c_2 + \delta V^{\pi_1}(x_1 + 1, 1).$$

The truth of this comparison depends on the state space X , so we break up this into 2 subcases.

Subcase 6.3.1: $(M_2 - 1) - (M_1 - x_1 - 1) \geq M_1$

By substituting in for V^{π_1} and Φ^{π_1} and then isolating c_0 , we can rewrite this inequality as

$$c_0 \geq \frac{-c_1 * \left(\frac{\delta^{M_1-x_1} - \delta^{M_2-x_2}}{1 - \delta^{M_1}} \right) - c_2 * \left(1 + \frac{\delta^{(\lfloor \frac{M_2-M_1+x_1}{M_1} \rfloor + 1)M_1-x_1} - \delta^{M_2-x_2}}{1 - \delta^{M_1 \lfloor \frac{M_2}{M_1} \rfloor}} \right)}{\left(1 + \frac{\delta^{M_1-x_1} - \delta^{M_2-x_2}}{1 - \delta^{M_1}} \right)} \quad (\text{B.10})$$

□

Subcase 6.3.2: $(M_2 - 1) - (M_1 - x_1 - 1) < M_1$

By substituting in for V^{π_1} and Φ^{π_1} and then isolating c_0 , we see that this is true for all $c_0 \geq 0$.

□

By looking across all cases, we have shown that if Equations B.4-B.10 are satisfied, then π_1 is the optimal decision policy. Each of these equations represents a lower bound on the visit cost c_0 , so if c_0 is large enough (i.e. above each of these thresholds for all possible values of X) then the optimal treatment policy is always the synchronized policy π_1 . Therefore, we can write \bar{c}_0 as the maximum of these thresholds. ■

B.2 Proof of Lemma 3.4.2

Recall that Lemma 3.4.2 states the following: *If $M_2 = kM_1$ then \bar{c}_0 is fully described by Equation B.7.*

To prove Lemma 3.4.2, we will look at each equation within Theorem 3.4.2 individually. Because we are considering \bar{c}_0 , we immediately know that we do not need to consider Equations B.4 and B.5 as they do not affect this threshold. For the other equations, we will utilize the definition of states within S_{π_1} as described in Equation 3.5. As a reminder, for all states in S_{π_1} , we have that $x_2 = x_1 + lM_1$, where l is any integer such that $0 \leq l < \lfloor \frac{M_2}{M_1} \rfloor - 1$. In our case, since $M_2 = kM_1$, we have that $0 \leq l < k - 1$.

We next show that when $M_2 = kM_1$, Equation B.6 is also irrelevant. This is because we never enter Case 3 (i.e., the states in Case 3 are not included in S_{π_1}). Assume instead that we would enter Case 3. Then, we be in a state such that

$$\begin{aligned}
& M_2 - x_2 < M_1 \\
\implies & kM_1 - x_2 < M_1 \\
\implies & kM_1 - x_1 - lM_1 < M_1 \\
\implies & kM_1 - M_1 - lM_1 < M_1 \\
\implies & (k - l)M_1 < 2M_1 \\
\implies & k - l < 2 \\
\implies & l \geq k - 1
\end{aligned}$$

which is a contradiction, since $l < k - 1$.

Similarly, we can show that Equation B.8 is irrelevant by contradiction. Assume instead that we can enter this subcase. Then, we would have that

$$\begin{aligned}
& (M_2 - x_2 - 1) - (M_1 - x_1 - 1) = M_1 - 1 \\
\implies & M_2 - x_2 - M_1 + x_1 = M_1 - 1 \\
\implies & kM_1 - x_1 - lM_1 - M_1 + x_1 = M_1 - 1 \\
\implies & (k - l)M_1 - M_1 = M_1 - 1 \\
\implies & (k - l - 2)M_1 = -1 \\
\implies & k - l = 2 - \frac{1}{M_1}
\end{aligned}$$

which is a contradiction, since $M_1 > 0$ and $l < k - 1 \implies k - l \geq 2$.

Similarly, we can show that Equation B.9 is irrelevant by contradiction. Assume instead that we can enter this subcase. Then, we would have that

$$\begin{aligned}
& M_2 - x_2 - 1 < M_1 - 1 \\
\implies & M_2 - x_2 < M_1 \\
\implies & kM_1 - x_2 < M_1 \\
\implies & kM_1 - x_1 - lM_1 < M_1 \\
\implies & (k - l - 1)M_1 - x_1 < 0 \\
\implies & k - 1 < 1 - \frac{x_1}{M_1}
\end{aligned}$$

which is a contradiction, since $x_1 < M_1 \implies 1 + \frac{x_1}{M_1} < 2 \implies k - l < 2 \implies l \geq k - 1$.

Similarly, we can show that Equation B.10 is irrelevant by contradiction. Assume instead that we can enter this subcase. Then, we would have that

$$\begin{aligned}
 & M_1 - x_1 \geq M_2 - x_2 \\
 \implies & M_1 - x_1 \geq kM_1 - x_1 - lM_1 \\
 \implies & (k - l - 1)M_1 \leq 0 \\
 \implies & k - l \leq 1 \\
 \implies & l \geq k - 1
 \end{aligned}$$

which is a contradiction.

Thus, we have shown that if $M_2 = kM_1$, then the \bar{c}_0 threshold is simply Equation B.7.

APPENDIX C

Supporting Material for Chapter 4

C.1 Organ Procurement and Transplantation Network (OPTN) Eligible Death Criteria

Source: https://optn.transplant.hrsa.gov/media/1200/optn_policies.pdf (Accessed August 08, 2021)

For reporting purposes of performance assessments, an eligible death for deceased organ donation is defined as the death of a patient who meets all the following characteristics:

- Is 75 years old or less
- Is legally declared dead by neurologic criteria according to state or local law
- Has body weight of 5 kg or greater
- Has a body mass index (BMI) of 50 kg/m² or less
- Has at least one kidney, liver, heart or lung that is deemed to meet the eligible data definition as defined below:
 - The kidney would initially meet the eligible data definition unless the donor meets any of the following criteria:
 - * Greater than 70 years old
 - * Age 50-69 years with history of type 1 diabetes for more than 20 years
 - * Polycystic kidney disease
 - * Glomerulosclerosis greater than or equal to 20% by kidney biopsy
 - * Terminal serum creatinine greater than 4.0 mg/dL

- * Chronic renal failure
 - * No urine output for 24 hours or longer
- The liver would initially meet the eligible data definition unless the donor meets any of the following criteria:
- * Cirrhosis
 - * Terminal total bilirubin greater than or equal to 4 mg/dL
 - * Portal hypertension
 - * Macrosteatosis greater than or equal to 50% or fibrosis greater than or equal to stage II
 - * Fulminant hepatic failure
 - * Terminal AST/ALT greater than 700 U/L
- The heart would initially meet the eligible data definition unless the donor meets any of the following criteria:
- * Greater than 60 years old
 - * 45 years old or older with a history of 10 or more years of HTN or 10 or more years of type 1 diabetes
 - * History of coronary artery bypass graft (CABG)
 - * History of coronary stent/intervention
 - * Current or past medical history of myocardial infarction (MI)
 - * Severe vessel diagnosis as supported by cardiac catheterization (that is more than 50 percent occlusion or 2+ vessel disease)
 - * Acute myocarditis or endocarditis, or both
 - * Heart failure due to cardiomyopathy
 - * Internal defibrillator or pacemaker
 - * Moderate to severe single valve or 2-valve disease documented by echo or cardiac catheterization, or previous valve repair
 - * Serial echo results showing severe global hypokinesis
 - * Myxoma
 - * Congenital defects (surgically corrected or not)
- The lung would initially meet the eligible data definition unless the donor meets any of the following criteria:
- * Greater than 65 years old

- * Diagnosed with COPD
- * Terminal PaO₂/FiO₂ less than 250 mmHg
- * Asthma (with daily prescription)
- * Asthma is the cause of death
- * Pulmonary fibrosis
- * Previous lobectomy Multiple blebs documented on computed axial tomography (CAT) scan
- * Pneumonia as indicated on computed tomography (CT), X-ray, bronchoscopy, or cultures
- * Bilateral severe pulmonary contusions as per CT

If a deceased patient meets the above criteria they would be classified as an eligible death unless the donor meets any of the following criteria:

- The donor goes to the operating room with intent to recover organs for transplant and all organs are deemed not medically suitable for transplant
- The donor exhibits any of the following active infections (with a specific diagnosis):
 - Bacterial: tuberculosis, gangrenous bowel or perforated bowel or intra-abdominal sepsis
 - Viral: HIV infection by serologic or molecular detection, rabies, reactive hepatitis B surface antigen, retroviral infections including viral encephalitis or meningitis, active herpes simplex, varicella zoster, or cytomegalovirus viremia or pneumonia, acute Epstein-Barr virus (mononucleosis), West Nile virus infection, or SARS. However, an HIV positive organ procured for transplantation into an HIV positive recipient at a transplant hospital that meets the requirements in Policy 15.7: Open Variance for the Recovery and Transplantation of Organs from HIV Positive Donors would still meet the requirements of an eligible death, according to the OPTN Final Rule.
 - Fungal: active infection with cryptococcus, aspergillus, histoplasma, coccidioides, active candidemia or invasive yeast infection
 - Parasites: active infection with trypanosoma cruzi (Chagas'), Leishmania, strongyloides, or malaria (plasmodium sp.)
 - Prion: Creutzfeldt-Jacob disease

The following are general exclusions:

- Aplastic anemia, agranulocytosis
- Current malignant neoplasms, except non-melanoma skin cancers such as basal cell and squamous cell cancer and primary CNS tumors without evident metastatic disease
- Previous malignant neoplasms with current evident metastatic disease
- A history of melanoma
- Hematologic malignancies: leukemia, Hodgkin's disease, lymphoma, multiple myeloma
- Active fungal, parasitic, viral, or bacterial meningitis or encephalitis
- No discernible cause of death




Figure C.1: Hazard ratios calculated via Cox regression regarding the association of organ procurement organization (OPO) with graft survival after controlling for donor eligibility, as well as recipient age, ethnicity, sex, and BMI. OPO names have been de-identified using a randomly assigned three-digit code. OPO 120 was randomly selected as the baseline OPO.

BIBLIOGRAPHY

- [1] Alberto Abadie and Sebastien Gay. The impact of presumed consent legislation on cadaveric organ donation: A cross-country study. *Journal of Health Economics*, 25(4):599–620, 2006.
- [2] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *Journal of Machine Learning Research*, volume 23, pages 39–40, 2012.
- [3] Vishal Ahuja and John R. Birge. Response-adaptive designs for clinical trials: Simultaneous learning from multiple patients. *European Journal of Operational Research*, 248(2):619–633, 2016.
- [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, Oct 2004.
- [5] Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. In *Journal of Machine Learning Research*, volume 3, pages 397–422, 2003.
- [6] David A. Axelrod, Mark A. Schnitzler, Huiling Xiao, William Irish, Elizabeth Tuttle-Newhall, Su Hsin Chang, Bertram L. Kasiske, Tarek Alhamad, and Krista L. Lentine. An economic assessment of contemporary kidney transplant practice. *American Journal of Transplantation*, 18(5):1168–1176, May 2018.
- [7] Turgay Ayer, Oguzhan Alagoz, and Natasha K. Stout. OR forum-A POMDP approach to personalize mammography screening decisions. *Operations Research*, 60(5):1019–1034, Sep 2012.
- [8] Turgay Ayer, Can Zhang, Anthony Bonifonte, Anne C. Spaulding, and Jagpreet Chhatwal. Prioritizing hepatitis C treatment in U.S. Prisons. *Operations Research*, 67(3):853–873, Mar 2019.
- [9] Sakine Batun, Andrew J. Schaefer, Atul Bhandari, and Mark S. Roberts. Optimal Liver Acceptance for Risk-Sensitive Patients. *Service Science*, 10(3):320–333, Sep 2018.
- [10] Karina Berg, Terje R. Pedersen, Leiv Sandvik, and Ragnheiour Bragadóttir. Comparison of Ranibizumab and Bevacizumab for Neovascular Age-Related Macular Degeneration According to LUCAS Treat-and-Extend Protocol. *Ophthalmology*, 122(1):146–152, Jan 2015.

- [11] Hrvoje Bogunovic, Michael D. Abramoff, Li Zhang, and Milan Sonka. Prediction of treatment response from retinal OCT in patients with exudative age-related macular degeneration. In *Ophthalmic Medical Image Analysis First International Workshop*, pages 129–136. Boston, MA, 2017.
- [12] R. Scott Braithwaite, John Concato, Chung Chou Chang, Mark S. Roberts, and Amy C. Justice. A framework for tailoring clinical guidelines to comorbidity at the point of care, nov 2007.
- [13] Gary C. Brown, Melissa M. Brown, Sara B. Rapuano, and David Boyer. A Cost-Benefit Analysis of VEGF-Inhibitor Therapy for Neovascular Age-Related Macular Degeneration in the United States. *American Journal of Ophthalmology*, 223:405–429, Mar 2021.
- [14] Asuncion Sancho Calabuig, Eva Gavela Martínez, Julia Kanter Berga, Irina Sanchis, Ana Avila Bernabeu, Josep Francesc Crespo Albiach, and Luis Manuel Pallardo Mateu. Is the Donor Age a Real Limitation to the Use of Very Extended Criteria Donors in Kidney Transplantation? *Transplantation*, 102(Supplement 7):S467, Jul 2018.
- [15] Andrew M. Cameron, R. Mark Ghobrial, Hasan Yersiz, Douglas G. Farmer, Gerald S. Lipschutz, Sherilyn A. Gordon, Michael Zimmerman, Johnny Hong, Thomas E. Collins, Jeffery Gornbein, Farin Amersi, Michael Weaver, Carlos Cao, Tony Chen, Jonathan R. Hiatt, and Ronald W. Busuttil. Optimal utilization of donor grafts with extended criteria: A single-center experience in over 1000 liver transplants. *Annals of Surgery*, 243(6):748–753, Jun 2006.
- [16] Centers for Disease Control. Common Eye Disorders — Basics — VHI — CDC, 2019.
- [17] Centers for Medicare and Medicaid Services. Physician Fee Schedule, Aug 2021.
- [18] Qiushi Chen, Turgay Ayer, and Jagpreet Chhatwal. Optimal M-switch surveillance policies for liver cancer in a hepatitis C-infected population. *Operations Research*, 66(3):673–696, May 2018.
- [19] K. P. Croome, D. D. Lee, A. P. Keaveny, and C. B. Taner. Noneligible Donors as a Strategy to Decrease the Organ Shortage. *American Journal of Transplantation*, 17(6):1649–1655, Jun 2017.
- [20] Tinglong Dai and Sridhar Tayur. Healthcare operations management: A snapshot of emerging research. *Manufacturing and Service Operations Management*, 22(5):869–887, 2021.
- [21] Luke J. DeRoos, Wesley J. Marrero, Elliot B. Tapper, Christopher J. Sonnenday, Mariel S. Lavieri, David W. Hutton, and Neehar D. Parikh. Estimated Association Between Organ Availability and Presumed Consent in Solid Organ Transplant. *JAMA network open*, 2(10):e1912431, Oct 2019.
- [22] Sascha Fauser, Viktoria Schwabecker, and Philipp S. Muether. Suppression of intraocular vascular endothelial growth factor during aflibercept treatment of age-related macular degeneration. *American Journal of Ophthalmology*, 158(3):532–536, 2014.

- [23] K. Bailey Freund, Jean François Korobelnik, Robert Devenyi, Carsten Framme, John Galic, Edward Herbert, Hans Hoerauf, Paolo Lanzetta, Stephan Michels, Paul Mitchell, Jordi Monés, Carl Regillo, Ramin Tadayoni, James Talks, and Sebastian Wolf. Treat-And-Extend Regimens with Anti-Vegf Agents in Retinal Diseases. *Retina*, 35(8):1489–1506, Aug 2015.
- [24] A. Gharbi and J. P. Kenné. Maintenance scheduling and production control of multiple-machine manufacturing systems. In *Computers and Industrial Engineering*, volume 48, pages 693–707. Pergamon, Jun 2005.
- [25] Audrey Giocanti-Auregan, Ramin Tadayoni, Typhaine Grenet, Franck Fajnkuchen, Sylvia Nghiem-Buffet, Corinne Delahaye-Mazza, Gabriel Quentel, and Salomon Y. Cohen. Estimation of the need for bilateral intravitreal anti-VEGF injections in clinical practice. *BMC ophthalmology*, 16(1), Aug 2016.
- [26] J C Gittins. Bandit Processes and Dynamic Allocation Indices. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(2):148–164, 1979.
- [27] Nicolas Goldaracena, J. Michael Cullen, Dong Sik Kim, Burcin Ekser, and Karim J. Halazun. Expanding the donor pool for liver transplantation with marginal donors, Oct 2020.
- [28] D. Goldberg, M. J. Kallan, L. Fu, M. Ciccarone, J. Ramirez, P. Rosenberg, J. Arnold, G. Segal, K. P. Moritsugu, H. Nathan, R. Hasz, and P. L. Abt. Changing Metrics of Organ Procurement Organization Performance in Order to Increase Organ Donation Rates in the United States. *American Journal of Transplantation*, 17(12):3183–3192, Dec 2017.
- [29] D. S. Goldberg, B. French, P. L. Abt, and R. K. Gilroy. Increasing the Number of Organ Transplants in the United States by Optimizing Donor Authorization Rates. *American Journal of Transplantation*, 15(8):2117–2125, Aug 2015.
- [30] David S. Goldberg, Brianna Doby, and Raymond Lynch. Addressing critiques of the proposed CMS metric of organ procurement organ performance: More data isn’t better. *Transplantation*, 104(8):1662–1667, Aug 2020.
- [31] The CATT Research Group. Ranibizumab and Bevacizumab for Neovascular Age-Related Macular Degeneration. <http://dx.doi.org/10.1056/NEJMoa1102673>, 364(20):1897–1908, 2011.
- [32] Omesh P. Gupta, Gary Shienbaum, Avni H. Patel, Christopher Fecarotta, Richard S. Kaiser, and Carl D. Regillo. A treat and extend regimen using ranibizumab for neovascular age-related macular degeneration: Clinical and economic impact. *Ophthalmology*, 117(11):2134–2140, 2010.
- [33] Ali Hajjar and Oguzhan Alagoz. Personalized Disease Screening Decisions Considering a Chronic Condition. <https://doi.org/10.1287/mnsc.2022.4336>, Mar 2022.
- [34] A. Hart, J. M. Smith, M. A. Skeans, S. K. Gustafson, A. R. Wilk, S. Castro, J. Foutz, J. L. Wainright, J. J. Snyder, B. L. Kasiske, and A. K. Israni. OPTN/SRTR 2018 Annual Data Report: Kidney. *American Journal of Transplantation*, 20(s1):20–130, Jan 2020.

- [35] IBM Watson Health. IBM Micromedex RED BOOK, Aug 2021.
- [36] Health Resources and Services Organization. 2012 National Survey of Organ Donation Attitudes and Behaviors. *National Survey of Organ Donation Attitudes and Behaviors*, 20004(202), 2013.
- [37] Jeffrey S. Heier et al. Efficacy, durability, and safety of intravitreal faricimab up to every 16 weeks for neovascular age-related macular degeneration (TENAYA and LUCERNE): two randomised, double-masked, phase 3, non-inferiority trials. *The Lancet*, 399(10326):729–740, feb 2022.
- [38] Jonathan E Helm, Mariel S Lavieri, Mark P Van Oyen, Joshua D Stein, and David C Musch. Dynamic forecasting and control algorithms of glaucoma progression for clinician decision support. *Operations Research*, 63(5):979–999, 2015.
- [39] Robert Hörster, Tina Ristau, Srinivas R. Sadda, and Sandra Liakopoulos. Individual recurrence intervals after anti-VEGF therapy for age-related macular degeneration. *Graefes’ Archive for Clinical and Experimental Ophthalmology*, 249(5):645–652, May 2011.
- [40] A. K. Israni, D. Zaun, N. Hadley, J. D. Rosendale, C. Schaffhausen, W. McKinney, J. J. Snyder, and B. L. Kasiske. OPTN/SRTR 2018 Annual Data Report: Deceased Organ Donation. *American Journal of Transplantation*, 20(s1):509–541, Jan 2020.
- [41] A. K. Israni, D. Zaun, J. D. Rosendale, C. Schaffhausen, J. J. Snyder, and B. L. Kasiske. OPTN/SRTR 2017 Annual Data Report: Deceased Organ Donation. *American journal of transplantation : official journal of the American Society of Transplantation and the American Society of Transplant Surgeons*, 19:485–516, Feb 2019.
- [42] Irfan Kabiruddin Jeeva, Sidra Masud, M. A.Rehman Siddiqui, and Hadees Murad Fahaad. Safety of simultaneous bilateral intravitreal versus unilateral anti-vasculo-endothelial growth factors injection in an operating room setting. *Pakistan journal of medical sciences*, 38(8):2324–2330, Nov 2022.
- [43] Leslie Pack Kaelbling. *Learning in Embedded Systems*. MIT Press, 2019.
- [44] Hossein Kamalzadeh, Vishal Ahuja, Michael Hahsler, and Michael E. Bowen. An Analytics-Driven Approach for Optimal Individualized Diabetes Screening. *Production and Operations Management*, 2021.
- [45] Jennifer J. Kang-Mieler, Kayla M. Rudeen, Wenqiang Liu, and William F. Mieler. Advances in ocular drug delivery systems, Feb 2020.
- [46] Steven M. Kawut, Alexander Reyentovich, Jessie S. Wilt, Roberto Anzeck, David J. Lederer, Mitchell K. O’Shea, Joshua R. Sonett, and Selim M. Arcasoy. Outcomes of extended donor lung recipients after lung transplantation. *Transplantation*, 79(3):310–316, Feb 2005.
- [47] I. Kennedy, R. A. Sells, A. S. Daar, R. D. Guttman, R. Hoffenberg, M. Lock, J. Radcliffe-Richards, and N. Tilney. The case for ‘presumed consent’ in organ donation, May 1998.

- [48] WR Kim, JR Lake, JM Smith, DP Schladt, MA Skeans, and SM Noreen. OPTN/SRTR 2017 Annual Data Report: Liver. *American Journal of Transplantation*, 19:184–283, Feb 2019.
- [49] Gina Koch, Bonnie J. Wakefield, and Douglas S. Wakefield. Barriers and Facilitators to Managing Multiple Chronic Conditions: A Systematic Literature Review. *Western Journal of Nursing Research*, 37(4):498–516, Apr 2015.
- [50] Sakhee Kotecha, Jamie Hobson, Jeremy Fuller, Eldho Paul, Bronwyn J. Levvey, Helen Whitford, Miranda Paraskeva, David McGiffin, Gregory I. Snell, and Glen P. Westall. Continued Successful Evolution of Extended Criteria Donor Lungs for Transplantation. *Annals of Thoracic Surgery*, 104(5):1702–1709, Nov 2017.
- [51] Paolo Lanzetta, Anat Loewenstein, and The Vision Academy Steering Committee. Fundamental principles of an anti-VEGF treatment regimen: optimal application of intravitreal anti-vascular endothelial growth factor therapy of macular diseases. *Graefe's Archive for Clinical and Experimental Ophthalmology*, 255(7):1259, Jul 2017.
- [52] Elliot Lee, Mariel S Lavieri, and Michael Volk. Optimal screening for hepatocellular carcinoma: A restless bandit model. *Manufacturing and Service Operations Management*, 21(1):198–212, 2019.
- [53] SJ Loeb, J Penrod, S Falkenstern Western Journal of . . . , and undefined 2003. Supporting older adults living with multiple chronic conditions. *journals.sagepub.com*, 25(1):8–29, Feb 2003.
- [54] RD Luce. Individual Choice Behavior: A theoretical analysis, New York, NY: John Willey and Sons, 1959.
- [55] R. Luskin and H. Nathan. Eligible Death Statistic: Not a True Measure of OPO Performance nor the Potential to Increase Transplantation, Aug 2015.
- [56] Sayeed K. Malek, Brandon J. Keys, Sanjaya Kumar, Edgar Milford, and Stefan G. Tullius. Racial and ethnic disparities in kidney transplantation. *Transplant International*, 24(5):419–424, May 2011.
- [57] Irmela Mantel, Angeliki Deli, Katia Iglesias, and Aude Ambresin. Prospective study evaluating the predictability of need for retreatment with intravitreal ranibizumab for age-related macular degeneration. *Graefe's Archive for Clinical and Experimental Ophthalmology*, 251(3):697–704, Mar 2013.
- [58] Hasina Maredia, Mary Grace Bowring, Allan B Massie, Sunjae Bae, Amber Kernodle, Shakirat Oyetunji, Christian Merlo, Robert S D Higgins, Dorry L Segev, and Errol L Bush. Better Understanding the Disparity Associated With Black Race in Heart Transplant Outcomes: A National Registry Analysis. *Circulation. Heart failure*, 14(2):e006107, Feb 2021.
- [59] J. E. Mason, B. T. Denton, N. D. Shah, and S. A. Smith. Optimizing the simultaneous management of blood pressure and cholesterol for type 2 diabetes patients. *European Journal of Operational Research*, 233(3):727–738, mar 2014.

- [60] H. U. Meier-Kriesche, F. K. Port, A. O. Ojo, S. M. Rudich, J. A. Hanson, D. M. Cibrik, A. B. Leichtman, and B. Kaplan. Effect of waiting time on renal transplant outcome. *Kidney International*, 58(3):1311–1317, Sep 2000.
- [61] Maria Messina, Davide Diena, Sergio Dellepiane, Gabriella Guzzo, Luca Lo Sardo, Fabrizio Fop, Giuseppe P. Segoloni, Antonio Amoroso, Paola Magistroni, and Luigi Biancone. Long-term outcomes and discard rate of kidneys by decade of extended criteria donor age. *Clinical Journal of the American Society of Nephrology*, 12(2):323–331, Feb 2017.
- [62] Richard Moore, Roy J. Thomas, and Chris Jones. Organ Donation in Wales: Time to Reflect, Dec 2018.
- [63] Gur Mosheiov. Multi-Machine Scheduling With Linear Deterioration. *INFOR: Information Systems and Operational Research*, 36(4):205–214, 1998.
- [64] Beat Müllhaupt, Dimitrios Dimitroulis, J. Tilman Gerlach, and Pierre Alain Clavien. Hot topics in liver transplantation: Organ allocation - extended criteria donor - living donor liver transplantation, Jan 2008.
- [65] Diana M Negoescu, Kostas Bimpikis, Margaret L Brandeau, and Dan A Iancu. Dynamic learning of patient response types: An application to treating chronic diseases. *Management Science*, 64(8):3469–3488, 2018.
- [66] Sweet Ping Ng, Temitayo Ajayi, Andrew J. Schaefer, Courtney Pollard, Houda Bahig, Adam S. Garden, David I. Rosenthal, G. Brandon Gunn, Steven J. Frank, Jack Phan, William H. Morrison, Jason M. Johnson, Abdallah S.R. Mohamed, Erich M. Sturgis, and Clifton D. Fuller. Surveillance imaging for patients with head and neck cancer treated with definitive radiotherapy: A partially observed Markov decision process model. *Cancer*, 126(4):749–756, Feb 2020.
- [67] Jane Noyes, Leah McLaughlin, Karen Morgan, Philip Walton, Rebecca Curtis, Susanna Madden, Abigail Roberts, and Michael Stephens. Short-term impact of introducing a soft opt-out organ donation system in Wales: Before and after study. *BMJ Open*, 9(4), Apr 2019.
- [68] The Office of the White House. Advancing American Kidney Health, Exec. Order No. 13879, 2019.
- [69] Anand K. Parekh, Richard A. Goodman, Catherine Gordon, and Howard K. Koh. Managing Multiple Chronic Conditions: A Strategic Framework for Improving Health Outcomes and Quality of Life. <http://dx.doi.org/10.1177/003335491112600403>, 126(4):460–471, Jul 2011.
- [70] Neehar D. Parikh, David Hutton, Wesley Marrero, Kunal Sanghani, Yongcai Xu, and Mariel Lavieri. Projections in donor organs available for liver transplantation in the United States: 2014-2025. *Liver Transplantation*, 21(6):855–863, jun 2015.
- [71] John F. Payne, Charles C. Wykoff, W. Lloyd Clark, Beau B. Bruce, David S. Boyer, and David M. Brown. Randomized Trial of Treat and Extend Ranibizumab with and without Navigated Laser for Diabetic Macular Edema: TREX-DME 1 Year Outcomes. *Ophthalmology*, 124(1):74–81, Jan 2017.

- [72] Jenny Ploeg, Nancy Matthew-Maich, Kimberly Fraser, Sinéad Dufour, Carrie McAiney, Sharon Kaasalainen, Maureen Markle-Reid, Ross Upshur, Laura Cleghorn, and Anna Emili. Managing multiple chronic conditions in the community: a Canadian qualitative study of the experiences of older adults, family caregivers and healthcare providers. *BMC Geriatrics*, 17(1):1–15, Jan 2017.
- [73] Jenny Ploeg, Marie Lee Yous, Kimberly Fraser, Sinéad Dufour, Lisa Garland Baird, Sharon Kaasalainen, Carrie McAiney, and Maureen Markle-Reid. Healthcare providers’ experiences in supporting community-living older adults to manage multiple chronic conditions: A qualitative study. *BMC Geriatrics*, 19(1), Nov 2019.
- [74] Adam M. Pourcho, Jay Smith, Stephen J. Wisniewski, and Jacob L. Sellon. Intraarticular platelet-rich plasma injection in the treatment of knee osteoarthritis: Review and recommendations. *American Journal of Physical Medicine and Rehabilitation*, 93(11):S108–S121, 2014.
- [75] Martin L. Puterman. *Markov decision processes: Discrete stochastic dynamic programming*. Wiley Series in Probability and Statistics. John Wiley & Sons, Inc., Hoboken, NJ, USA, Apr 2008.
- [76] Abbas Rana, Angelika Gruessner, Vatche G. Agopian, Zain Khalpey, Irbaz B. Riaz, Bruce Kaplan, Karim J. Halazun, Ronald W. Busuttill, and Rainer W.G. Gruessner. Survival benefit of solid-organ transplant in the United States. *JAMA Surgery*, 150(3):252–259, Mar 2015.
- [77] Amber Rithalia, Catriona McDaid, Sara Suekarran, Lindsey Myers, and Amanda Sowden. Impact of presumed consent for organ donation on donation rates: A systematic review. *BMJ (Online)*, 338(7689):284–287, Jan 2009.
- [78] Markus Rohm, Volker Tresp, Michael Müller, Christoph Kern, Ilja Manakov, Maximilian Weiss, Dawn A. Sim, Siegfried Priglinger, Pearse A. Keane, and Karsten Kortuem. Predicting Visual Acuity by Using Machine Learning in Patients Treated for Neovascular Age-Related Macular Degeneration. *Ophthalmology*, 125(7):1028–1036, 2018.
- [79] Dronacharya Routh, Sanjay Sharma, C. S. Naidu, P. P. Rao, Anuj Kumar Sharma, and Priya Ranjan. Comparison of outcomes in ideal donor and extended criteria donor in deceased donor liver transplant: A prospective study. *International Journal of Surgery*, 12(8):774–777, Aug 2014.
- [80] Miguel Ruão, María Andreu-Fenoll, Rosa Dolz-Marco, and Roberto Gallego-Pinazo. Safety of bilateral same-day intravitreal injections of anti-vascular endothelial growth factor agents. *Clinical ophthalmology (Auckland, N.Z.)*, 11:299–302, Feb 2017.
- [81] Diego Ruiz-Hernández, Jesús M Pinar-Pérez, and David Delgado-Gómez. Multi-machine preventive maintenance scheduling with imperfect interventions: A restless bandit approach. *Computers & Operations Research*, 119:104927, 2020.
- [82] Sammy Saab, Satvir S. Saggi, Mizna Akbar, and Gina Choi. Presumed Consent: A Potential Tool for Countries Experiencing an Organ Donation Crisis. *Digestive Diseases and Sciences*, 64(5):1346–1355, May 2019.

- [83] Alireza Sabouri, Woonghee Tim Huh, and Steven M. Shechter. Screening strategies for patients on the kidney transplant waiting list. *Operations Research*, 65(5):1131–1146, Aug 2017.
- [84] Gregory 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.
- [85] Thomas Schlegl, Sebastian M. Waldstein, Hrvoje Bogunovic, Franz Endstraßer, Amir Sadeghipour, Ana Maria Philip, Dominika Podkowinski, Bianca S. Gerendas, Georg Langs, and Ursula Schmidt-Erfurth. Fully Automated Detection and Quantification of Macular Fluid in OCT Using Deep Learning. *Ophthalmology*, 125(4):549–558, 2018.
- [86] Ursula Schmidt-Erfurth, Hrvoje Bogunovic, Amir Sadeghipour, Thomas Schlegl, Georg Langs, Bianca S. Gerendas, Aaron Osborne, and Sebastian M. Waldstein. Machine Learning to Analyze the Prognostic Value of Current Imaging Biomarkers in Neovascular Age-Related Macular Degeneration. *Ophthalmology Retina*, 2(1):24–30, 2018.
- [87] 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, Jan 2008.
- [88] Ellen Sheehy, Suzanne L. Conrad, Lori E. Brigham, Richard Luskin, Phyllis Weber, Mark Eakin, Lawrence Schkade, and Lawrence Hunsicker. Estimating the number of potential organ donors in the United States, Aug 2003.
- [89] M. Reza Skandari and Steven M. Shechter. Patient-type bayes-adaptive treatment plans. *Operations Research*, 69(2):574–598, Feb 2021.
- [90] Christine Hahjin Song, Simon Stern, Mohana Giruparajah, Noam Berlin, and Gordon L. Sussman. Long-term efficacy of fixed-dose omalizumab for patients with severe chronic spontaneous urticaria. *Annals of Allergy, Asthma and Immunology*, 110(2):113–117, Feb 2013.
- [91] Lauren N. Steimle and Brian T. Denton. Markov decision processes for screening and treatment of chronic diseases. In *International Series in Operations Research and Management Science*, volume 248, pages 189–222. Springer New York LLC, 2017.
- [92] Glen B. Taksler, Melanie Keshner, Angela Fagerlin, Negin Hajizadeh, and R. Scott Braithwaite. Personalized estimates of benefit from preventive care guidelines: A proof of concept. *Annals of Internal Medicine*, 159(3):161–168, Aug 2013.
- [93] Rekha V. Thammana, Stuart J. Knechtle, Rene Romero, Thomas G. Heffron, Caroline T. Daniels, and Rachel E. Patzer. Racial and socioeconomic disparities in pediatric and young adult liver transplant outcomes. *Liver Transplantation*, 20(1):100–115, Jan 2014.
- [94] CJCH Watkins. *Learning from delayed rewards*. PhD thesis, Cambridge, 1989.

- [95] P. Whittle. Restless bandits: activity allocation in a changing world. *Journal of Applied Probability*, 25(A):287–298, 1988.
- [96] Thomas R. Wojda, Stanislaw P. Stawicki, Kathy P. Yandle, Maria Bleil, Jennifer Axelband, Rebecca Wilde-Onia, Peter G. Thomas, James Cipolla, William S. Hoff, and Jill Shultz. Keys to successful organ procurement: An experience-based review of clinical practices at a high-performing health-care organization. *International Journal of Critical Illness and Injury Science*, 7(2):91–100, apr 2017.
- [97] Michael Woodroffe. A one-armed bandit problem with a concomitant variable. *Journal of the American Statistical Association*, 74(368):799–806, 1979.
- [98] Charles C. Wykoff, William C. Ou, David M. Brown, Daniel E. Croft, Rui Wang, John F. Payne, W. Lloyd Clark, Nizar Saleh Abdelfattah, and Srini Vas R. Sadda. Randomized trial of treat-and-extend versus monthly dosing for neovascular age-related macular degeneration: 2-year results of the trex-amd study. *Ophthalmology Retina*, 1:314–321, 7 2017.
- [99] Alejandra Zúñiga-Fajuri. Increasing organ donation by presumed consent and allocation priority: Chile. *Bulletin of the World Health Organization*, 93(3):199–202, 2015.