OPTIMAL POLICIES FOR THE ACCEPTANCE OF LIVING- AND CADAVERIC-DONOR LIVERS

by

Oguzhan Alagoz

BS, Bilkent University, 1997

MS, Middle East Technical University, 2000

Submitted to the Graduate Faculty of

the School of Engineering in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2004

UNIVERSITY OF PITTSBURGH

SCHOOL OF ENGINEERING

This dissertation was presented

by

Oguzhan Alagoz

It was defended on

July 12, 2004

and approved by

Andrew J. Schaefer, Assistant Professor, School of Engineering

Matthew Bailey, Assistant Professor, School of Engineering

Cindy L. Bryce, Assistant Professor, School of Medicine

Lisa M. Maillart, Assistant Professor, Weatherhead School of Management

Mainak Mazumdar, Professor, School of Engineering

Mark S. Roberts, Associate Professor, School of Medicine

Dissertation Director: Andrew J. Schaefer, Assistant Professor, School of Engineering

OPTIMAL POLICIES FOR THE ACCEPTANCE OF LIVING- AND CADAVERIC-DONOR LIVERS

Oguzhan Alagoz, PhD

University of Pittsburgh, 2004

Transplantation is the only viable therapy for end-stage liver diseases (ESLD) such as hepatitis B. In the United States, patients with ESLD are placed on a waiting list. When organs become available, they are offered to the patients on this waiting list. This dissertation focuses on the decision problem faced by these patients: which offer to accept and which to refuse? This decision depends on two major components: the patient's current and future health, as well as the current and future prospect for organ offers. A recent analysis of liver transplant data indicates that 60% of all livers offered to patients for transplantation are refused.

This problem is formulated as a discrete-time Markov decision process (MDP). This dissertation analyzes three MDP models, each representing a different situation. The Living-Donor-Only Model considers the problem of optimal timing of living-donor liver transplantation, which is accomplished by removing an entire lobe of a living donor's liver and implanting it into the recipient. The Cadaveric-Donor-Only Model considers the problem of accepting/refusing a cadaveric liver offer when the patient is on the waiting list but has no available living donor. In this model, the effect of the waiting list is incorporated into the decision model implicitly through the probability of being offered a liver. The Living-and-Cadaveric-Donor Model is the most general model. This model combines the first two models, in that the patient is both listed on the waiting list and also has an available living donor. The patient can accept the cadaveric liver offer, decline the cadaveric liver offer and use the living-donor liver, or decline both and continue to wait.

This dissertation derives structural properties of all three models, including several sets of conditions that ensure the existence of intuitively structured policies such as control-limit policies. The computational experiments use clinical data, and show that the optimal policy is typically of control-limit type.

Keywords: Control-limit policy, Markov decision processes, medical decision making, organ transplantation, service operations.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

To my wonderful parents, Zeliha and Ilhami Alagoz

ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to my advisor and mentor, Professor Andrew Schaefer for directing this study and supporting me through my PhD education. He has always been available for me. Many thanks to him for his generous support, advice, flexibility, inspiration and motivation.

My sincere thanks and appreciation go to my committee member, Professor Lisa Maillart for her help in conducting this research and preparing this manuscript. Her technical and editorial advice were essential to the completion of this dissertation. I really enjoyed working with her throughout this study.

I would also like to thank my committee member, Professor Mark Roberts for his valuable comments and enthusiasm throughout this study. I am also indebted to the rest of the members of my dissertation committee, Matthew Bailey, Cindy Bryce and Mainak Mazumdar for their valuable suggestions and insights. I am grateful to my former advisors Professor Meral Azizoglu and Professor Bryan Norman for encouraging me to pursue an academic career.

I am also grateful to my friends in the Computational Optimization Lab who mentally supported me throughout my study. Among them, special thanks to Nan Kong, Jennifer Kreke and Steven Shecthter, for their valuable insights and comments about my research. Thanks to my undergraduate helpers Danielle Williams and Laura Henderson for their help in preparing this manuscript.

I thank the wonderful staff of the Industrial Engineering Department, Lisa Bopp, Richard Brown, Minerva Hubbard and Jim Segneff, for providing technical support throughout my study.

Lastly, I am forever indebted to my wonderful parents Zeliha and Ilhami, to my dear sister Bilgehan and to my dear brother Mehmet Sezgin. I would have never finished this dissertation without their endless love, encouragement and unconditional support.

1.0 INTRODUCTION

The life expectancy for Americans increased from 49 years in 1900 to an all-time high of 77.2 years in 2001 [35, 56, 57]. Although factors such as changes in life style contribute to this advance, it is mostly due to improved public health policy and improvements in medical care [119]. A report prepared by MEDTAP International shows that, without the advances in healthcare, in 2000 alone there would have been 470,000 more deaths, 2.3 million more disabled persons, and 206 million more hospital days in the U.S. [84].

The Centers for Medicare and Medicaid Services (CMS) reports that in 2001, health-care spending in the nation rose to $1.4 trillion, 14% of the gross domestic product (GDP), making the health-service industry the largest industry in the U.S. [58, 60, 109]. Projections show that by 2011 this industry will constitute approximately 17% of GDP [58]. These data suggest that any improvement in health-care service might have significant effects on the overall economy as well as life expectancy.

Medical decision making, which is required at every stage of medical intervention, is the process of identifying, evaluating, and solving problems in a health-care system. Optimization models are a useful tool in medical decision making. In the last few decades, optimization has been used to determine the optimal use of physical resources in delivering medical care and improving service quality. Several researchers [89, 111, 153] consider personnel scheduling problems, particularly nurse scheduling. Other examples of the optimal use of physical entities in medical decision making include emergency room scheduling [8, 34, 100], scheduling of patient admissions [82], and ambulance location [18, 19].

Recently, however, optimization has been used for the treatment of individual patients, using physiological knowledge of the individual patient. Optimal cancer therapy has been a subject of much recent research interest. Lee et al. [97] consider the optimal placement

of radioactive seeds for brachytherapy for prostate cancer. Preciado-Walters et al. [123] use optimization for direct-beam cancer treatment. Lyberatos and Abulesz [101] evaluate the effect of periodic injections as a cancer chemotherapy regimen using an optimization model. Gallagher and Lee [63] use mixed integer programming optimization models for generating treatment plans for brachytherapy. Yoo et al. [164] consider the problem of planning treatment for prostate implants. Zhang et al. [167] describe an optimization algorithm for gamma knife radiosurgery treatment planning. Korotaev et al. [93] consider the problem of optimizing the dosage of a medication to obtain the best therapeutic effect. They apply their methodology to the problem of dosage optimization for an antiseptic drug in the treatment of pyelonephritis. Falk et al. [53] use optimization to develop bounds on a trauma outcome function. Magni et al. [102] apply a Markov decision process model to decide the optimal timing of the prophylactic surgery in mild hereditary spherocytosis. Hauskrecht and Fraser [74] consider the problem of planning treatment of ischemic heart disease and solve it by using a partially observable Markov decision processes model. Other researchers [9, 10, 105, 106] focus on the optimal sequence of chemotherapy treatments.

This dissertation focuses on optimal timing of organ transplantation. In today's medical world, donated human organs available for transplant are one of the most scarce resources. Although the number of organ transplants performed has increased annually since the enactment of the National Organ Transplant Act of 1984, the number of patients waiting for an organ transplant and the number of deaths have also increased dramatically [85]. In 2002, nearly 24,900 Americans were transplanted with a kidney, liver, heart or other organ [149]. On the other hand, in 2002, there were nearly 80,483 patients on the organ waiting lists and more than 6,000 of them died while waiting for an organ [149]. Given this scarcity, it is not surprising that the allocation and efficient usage of organs are very critical and controversial topics.

Table 1.1 shows the liver-related data in the U.S. between 1996-2001 [149]. As presented in the table, the number of patients and the deaths while waiting for a liver have increased dramatically in the last five years. Unfortunately, also as shown in the table, each year around 10% of all donated livers are wasted and 13% of all patients who receive transplants

die within the first year. These data suggest that there is a need to improve the allocation and usage of available organs.

The liver is an appropriate organ for considering accept-reject decisions. First, none of the previous studies on optimal organ allocation [1, 40, 41, 77, 126, 142, 165, 166] consider liver transplantation. Second, the liver data reflect the trends in the overall organ data such as imbalance between the number of organs and the number of patients on the waiting list, dramatic increase in the number of patients waiting over time, and a significant amount of organs wasted [149]. Moreover, although there are more kidney transplants than liver transplants, there is no alternative therapy for end-stage liver diseases (ESLD) such as primary biliary cirrhosis and hepatitis B, such as dialysis for kidney patients. When Congress instructed the Institute of Medicine to study organ transplantation, it focused on liver transplantation because "much of the current debate has centered on the procurement and allocation of livers" [85].

Many approaches have been proposed to improve the discrepancy between the number of sick patients and the number of available organs by increasing the number of organ donors. For instance, tax credits for the families of organ donors would provide a financial incentive for donation. Recently, the American Medical Association considered a proposal that would pay donors for organs [23]. In Boston and Washington, D.C., a family member can improve a relative's priority for a transplant by becoming an organ donor [118]. The Wisconsin Senate recently approved a bill calling for tax deduction for organ donors to cover their expenses [114]. Transplant surgeons are now splitting livers and transplanting them into two separate patients [104].

Most liver donations are from cadaveric donors. However, an increasing source of donated livers is *living donors*, in which transplantation is accomplished by removing an entire lobe of the donor's liver and implanting it into the recipient [83]. The non-diseased liver has a unique regenerative ability, so that a donor's liver regains its previous size within two weeks [83]; the same process occurs in the recipient. Living donors are often related to the recipient. While donating part of a liver does entail real risk for the donor, to date there have been only two donor deaths due to living-donor liver donation in the United States [131], which is comparable to the mortality rate of other living-donor procedures such as kidneys [36].

Table 1.1: U.S. Liver Data between 1996-2001 [149]

	1996	1997	1998	1999	2000	2001
Patients Waiting	7,398	9,527	11,908	14,445	16,874	18,214
Deaths	1,003	1,199	1,437	1,850	1,787	2,003
Transplants	4,077	4,185	4,506	4,726	4,969	5,184
Deceased donors	4017	4101	4416	4485	4582	4665
Adult Recipients	3547	3613	3896	4040	4118	4175
Pediatric Recipients	470	488	520	445	464	490
Living donors	60	84	87	230	380	515
Adult Recipients	3	3	24	140	272	408
Pediatric Recipients	57	81	63	90	108	107
Wasted Organs[1]	427	402	458	428	496	444
Survival rate (Cadaveric Donors)[2]	85.7%	87.5%	87.2%	86.5%	88.3%	87.7%
Survival rate (Living Donors)[3]	89.4%	83.6%	83.3%	81.0%	89.7%	91.2%

[1]An organ is assumed to be wasted only if it is donated by a cadaveric recipient, is viable for transplantation, but is not used for transplantation
[2]Annual post-transplant survival rate for cadaveric organ recipients
[3]Annual post-transplant survival rate for living-donor organ recipients

In this study, the most general model considers the availability of a living donor as well as potential cadaveric liver offers.

Due to a shortage of cadaveric livers, living donors have recently been used in liver transplantation. In fact, living donors are practically the only source of organs in countries where the availability of cadaveric donors is severely restricted, such as Japan [115]. There are several potential advantages of using a living donor: the organ is usually of higher quality, there is no waiting time on the list for the patient, the *cold ischemia* time, the time that a liver stays outside the body, is essentially zero, and a complete pre-operative evaluation of the donated liver is possible, which may increase the success of the operation [73]. In addition to these benefits, the post-transplant quality of life is generally higher for living-donor recipients than cadaveric liver recipients [145]. Moreover, the time of the transplantation can be selected rather than dictated by a cadaveric donation [83]. This timing decision is the focus of Chapter 5.

As can be seen from Table 1.1, although the number of living donors is still a small portion of all transplanted organs, the rate of growth in living-donor transplants is much greater than that of cadaveric transplants. The number of living-donor liver transplants has grown by an order of magnitude from 1996 to 2001, and now represent over 10% of all transplants. Initially, living-donor transplantation was almost exclusively used in the pediatric population, where a parent of a child with ESLD would donate a lobe in order to save the child's life. However, as demonstrated in Table 1.1, a large component of the rise in living-donor transplants has been from adult to adult. Bolstered by the success of adult-to-adult living-donor transplantation from Japan, where there are social prohibitions to cadaveric transplantation [115], adult-to-adult living-donor transplantation has spread rapidly in the United States as waiting list death rates have risen, and waiting time for organs has skyrocketed, with a median waiting time of 770 days for transplant recipients in 1998.

Previous research has focused on the optimal allocation of a given set of organs to a group of patients. Much of the research for the optimal distribution of organs for transplantation has focused on global optimization systems considering the societal perspective rather than the patient perspective. These researchers seek to provide an optimal match between organs

and patients by considering factors such as expected quality-adjusted life years, one-year graft survival rates, and quality of the prospective match to find an optimal allocation of organs [165, 166]. However, such an optimization model must address the ethical issues that arise when one group of patients has a higher priority for organs than another. Furthermore, the global optimal policy for allocation of organs ignores the fact that patients make their own decisions in current practice. Therefore, the implementation of the policies resulting from such models is very difficult. Some researchers have considered the question of which patients should receive organs [107, 147]. Many factors must be considered, including differing moral viewpoints [66] and political issues such as states' rights [146].

1.1 PROBLEM STATEMENT

The purpose of this dissertation is to determine whether a patient should accept or decline a liver offer, given her current health state. That is, we seek a policy describing the health state/liver type combinations in which the liver transplantation should occur, and those combinations where waiting is the optimal action.

The patient, the transplant surgeon and/or the physician responsible for the care of the patient act as a single decision maker [83, 149], and we assume that their objective is to maximize the patient's total expected discounted reward. The consideration of other objectives is left for future research. Discounting future health effects is a common practice in medical research (see Chapter 3 for a more detailed discussion of this practice). The model presented in this dissertation is general enough to handle various reward function definitions. Possible definitions include total discounted expected life days and total discounted quality-adjusted life days (QALD) of the patient, a common measure in medical decision making research. The QALD measure is based on the assumption that the patient assigns a quality score between 0 and 1 to each health state [68]. A quality score of 1 corresponds to perfect health whereas a quality score of 0 corresponds to death.

There are several advantages to considering the patient's perspective rather than society's perspective. The resulting optimal policies are immediately implementable by any patient within the current system, without the approval of any government agency. This problem

also reflects an actual decision made by thousands of patients every year. Considering the number of deaths and the number of unsuccessful transplant operations, an optimal solution to this problem may result in better usage of the organs and save many lives.

At first glance, the answer to this question seems very straightforward: a very sick patient should perhaps always accept a high quality organ while a relatively healthy patient should perhaps always decline a poor organ, but the decision is often more complicated in practice. A high risk of mortality without a transplant might cause a very sick patient to accept a poor organ, even if the post-transplant survival rate is not very high. Similarly, uncertainty in the availability and the quality of future liver offers might cause a healthy patient to accept a lower quality organ. On the other hand, if a relatively healthy patient is offered a poor quality organ while she is likely to be offered higher quality organs in the near future, she might prefer declining the organ offered.

The decision depends on two major components: the patient's current and future health as well as the current and future prospects for organs. In both cases, the future is unknown and the probabilities of different outcomes are difficult to estimate.

This problem may appear to be equivalent to the classical secretary problem [25, 125], where a decision maker is sequentially presented with a set of N candidates, where candidate i is of quality q_i. At stage i, the decision maker can either stop with candidate i and receive a reward of q_i, or he may continue and see candidate $i+1$. If the decision maker rejects candidate i, she seeks employment elsewhere. If the first $N-1$ candidates are rejected, candidate N must be accepted. The objective is to maximize the probability of choosing the best candidate. This variant of the secretary problem has an optimal closed-form solution in which the decision maker rejects the first 36.8% of candidates and then chooses the first candidate thereafter who is the better than the best candidate in the first 36.8% of candidates [125]. If no such candidate appears, the last candidate is chosen.

The problem described in this dissertation differs in several respects from the classical secretary problem. First, the future availability of livers is uncertain. Patients do not wait in a single queue for livers. Rather, different livers may induce different queues, depending on other characteristics such as the geographical location of the donor and the blood type of the liver. For one potential liver, patient A may have a higher priority than patient

B. However, for another liver, patient B may have the higher priority. Second, the set of potential candidates is essentially infinite and unknown; the patient might get organ offers as long as she is sick. Third, the decision maker's health changes over time, therefore the reward is dynamic. The exact progression of health states is unknown and may depend on such factors as disease. Lastly, unlike the secretary problem, there are intermediate rewards, i.e. as the patient waits she accumulates reward. These differences make the problem far more complicated than the classical secretary problem, and a closed-form solution is not likely to exist.

1.2 CURRENT LIVER ALLOCATION SYSTEM

This section describes the current liver allocation system. Basic knowledge of this system is necessary to understand the decision problem faced by the ESLD patients and the development of the decision model. The United Network for Organ Sharing (UNOS) is responsible for managing the national organ donation and allocation system. The UNOS Board of Directors approved for implementation the new liver allocation procedure as of February 28, 2002 [149]. It is important to note that, according to the General Accounting Office, in the last six years the policy has been changed four times [117, 149]. The multiple changes in policy over a short time period is evidence of the ever-changing opinions surrounding the optimal allocation of livers. Although the new policy is anticipated to "better identify urgent patients and reduce deaths among patients awaiting liver transplants" [149], anecdotal evidence suggests that there is some question among the transplant community as to whether the new allocation rules are satisfactory [64, 144].

UNOS manages the organ donation and procurement via Organ Procurement Organizations (OPOs) which are non-profit agencies responsible for approaching families about donation, evaluating the medical suitability of potential donors, coordinating the recovery, preservation, and transportation of organs donated for transplantation, and educating the public about the critical need for organ donation. There are currently 60 OPOs that operate in designated service areas; these service areas may cover multiple states, a single state, or just parts of a state [149]. The national UNOS membership is also divided into 11 geo-

graphic regions, each consisting of several OPOs. This regional structure was developed to facilitate organ allocation and to provide individuals with the opportunity to identify concerns regarding organ procurement, allocation, and transplantation that are unique to their particular geographic area [149]. Stahl et al. [141] consider the problem of optimal sizing and configuration of these regions and OPOs, and they provide a framework to find a set of regions that maximizes intra-regional transplants.

UNOS has different procedures for adult and for pediatric patients. In this study, we only consider adult patients; therefore, we will describe only the adult liver allocation procedure. UNOS maintains a patient waiting list that is used to determine the transplant candidates among the patients. Under the current policy, when a liver becomes available, the following factors are considered for its allocation: liver and patient OPO, liver and patient region, medical urgency of the patient, patient points, and patient waiting time.

The medical urgency of the adult liver patients is represented by UNOS Status 1 and MELD scores. According to the new UNOS policy, a patient listed as Status 1 "has fulminant liver failure with a life expectancy without a liver transplant of less than 7 days" [149]. Patients who do not qualify for classification as Status 1 do not receive a status level. Rather, these patients will be assigned a "probability of pre-transplant death derived from a mortality risk score" calculated by the Model for End Stage Liver Disease (MELD) scoring system [149]. The MELD score, which is a continuous function of total bilirubin, creatinine and prothrombin time, indicates the status of the liver disease and is a risk-prediction model first introduced by Malinchoc et al. to assess the short-term prognosis of patients with liver cirrhosis [103, 162]. Wiesner et al. [162] develop the following formula for computing MELD scores:

$$\text{MELD Score} = 10 \times [0.957 \times ln(\text{creatinine mg/DL}) + 0.378 \times ln(\text{bilirubin mg/DL})$$

$$+1.120 \times ln(\text{INR}) + 0.643 \times I_c]$$

where INR, international normalized ratio, is computed by dividing prothrombin time (PT) of the patient by a normal PT value, and I_c is an indicator variable that shows the cause of cirrhosis, i.e., it is equal to 1 if the disease is alcohol or cholestatic related and it is equal to 0 if the disease is related to other etiologies. As Wiesner et al. [162] note, the etiology

(cause) of disease is removed from the formula by UNOS. In addition to this, UNOS makes several modifications to the formula such as any lab value less than 1 mg/DL is set to 1 mg/DL, any creatinine level above 4 mg/DL is set to 4 mg/DL and the resulting MELD score is rounded to the closest integer[149]. By introducing these changes, UNOS restricts the range of MELD scores to be between 6 and 40, where a value of 6 corresponds to the best possible patient health and 40 to the worst.

Kamath et al. [88] introduce the MELD system to more accurately measure the liver disease severity and to better predict which patients are at risk of dying. However, there are concerns about the accuracy of the MELD system. First, there were some biases in the data used to develop the model. For instance, the data available to the researchers were mostly based on patients with advanced liver disease [103]. Furthermore, the MELD system was validated on the patients suffering from cirrhosis [162], therefore it is possible that the MELD system does not accurately measure the disease progression for other diseases, such as acute liver diseases. Moreover, as stated, although they presented data to indicate that the consideration of patient age, sex, and body mass is unlikely to be clinically significant, it is possible that other factors, such as a more direct measurement of renal function (iothalamate clearance), may improve the accuracy of the model [88]. For instance, patients who have hepatocellular carcinoma or metabolic diseases might have good renal function, but must wait because they are assigned a lower MELD score [99]. Furthermore, the MELD system was validated on only three laboratory values: creatinine and bilirubin levels, and prothrombin time. Thus, it is possible that the MELD system does not accurately consider patients with liver cancer because they would score as if they were healthy [64]. Consequently, relying mainly on laboratory results may not be the best solution for all patients [52].

Patients are stratified within Status 1 and each MELD score using patient "points" and waiting time. Patient points are assigned based on the compatibility of their blood type with the donor's blood type. For Status 1 patients, candidates with an exact blood type match receive 10 points; candidates with a compatible, though not identical, blood type receive 5 points; and a candidate whose blood type is incompatible receives 0 points. As an exception, though type O and type A_2 (a less common variant of blood type A) are incompatible, patients of type O receive 5 points for being willing to accept a type A_2 liver.

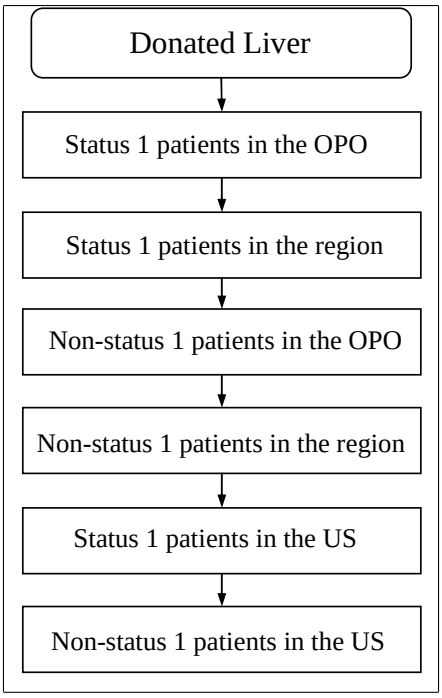

Figure 1.1: Current Liver Allocation System

For non-Status 1 patients with the same MELD score, a liver is offered to patients with an exact blood type match first, compatible patients second, and incompatible patients last. If there are several patients having the same blood type compatibility and MELD scores, the ties are broken with patient waiting time. The waiting time for a Status 1 patient is calculated only from the date when that patient was listed as Status 1. Points are assigned to each patient based on the following strategy: "Ten points will be accrued by the patient waiting for the longest period for a liver transplant and proportionately fewer points will be accrued by those patients with shorter tenure" [149]. For MELD patients, waiting time is calculated as the time accrued by the patient at or above her current score level from the date that she was listed as a candidate for liver transplantation.

Figure 1.1 shows a schematic representation of the liver allocation system. In summary, the current liver allocation system works as follows: every liver available for transplant is first offered to those Status 1 patients located within the harvesting OPO. When more than one Status 1 patient exists, the liver is offered to those patients in descending point order where the patient with the highest number of points receives the highest priority. If there are

no suitable Status 1 matches within the harvesting OPO, the liver is then offered to Status 1 patients within the harvesting region. If a match still has not been found, the liver is offered to all non-Status 1 patients in the harvesting OPO in descending order of MELD score. The search is again broadened to the harvesting region if no suitable match has been found. If no suitable match exists in the harvesting region, then the liver is offered nationally to Status 1 patients followed by all other patients in descending order of MELD scores.

UNOS maintains that the final decision to accept or decline a liver "will remain the prerogative of the transplant surgeon and/or physician responsible for the care of that patient" [85]. The surgeon and/or the physician have very limited time, namely one hour, to make their decision [149], because the acceptable range for cold ischemia time is very limited. The Scientific Registry of Transplant Recipients states that the acceptable cold ischemia time limit for a liver is 12 to 18 hours [116], whereas the Center for Organ Recovery and Education gives the maximum limit as 18 to 24 hours [59]. Furthermore, as the Institute of Medicine points out, there is evidence that the quality of the organ decreases as cold ischemia time increases [85]. In the event that a liver is declined, it is then offered to another patient in accordance with the above-described policy. The patient who declines the organ will not be penalized and will have access to future livers. Organs are frequently declined due to low quality of the liver. For example, the donor may have had health problems that could have damaged the organ or may be much older than the potential recipient, making the organ undesirable [154].

In current practice, organs are frequently declined. As we demonstrate in Section 4.4, 60% of all liver offers are declined. Howard [79] reports that 45% of livers are rejected by the first surgeon to whom they are offered. Although a lower percentage may hold for other organs, the problem addressed in this dissertation is faced, on average, by approximately 100 patients every day in the United States alone and is not expected to change in the near future [149].

The remainder of this dissertation is organized as follows: Chapter 2 discusses the literature related to Markov decision processes (MDP), MDP applications in medical decision making, simulation models of liver allocation system, and the organ allocation problem. Chapter 3 formulates the problem faced by the decision maker, namely, whether or not to

accept a liver offered for transplantation. Chapter 4 describes the data sources that are used in the computational study and presents three supplementary models that are input to the decision models in the subsequent chapters. Chapter 5 describes the Living-Donor-Only Model, in which the liver is assumed to be available at all points in time and its characteristics are known and constant. Chapter 6 describes the Cadaveric-Donor-Only Model, in which multiple patients compete for multiple livers. The other patients competing for livers are implicitly modeled through the probability of being offered livers in the future. Chapter 7 considers the combined model, in which the patient is assumed to have an available living donor and is listed on the waiting list. Chapter 8 summarizes the conclusions drawn based on these modes and gives possible future research directions.

2.0 LITERATURE REVIEW

In this chapter, we discuss the literature related to our study. In Section 2.1 we give a brief introduction to MDP models, describe some different applications of MDPs, and summarize several solution techniques. In Section 2.2 we summarize previous studies that apply MDPs to medical decision making problems. Section 2.3 discusses previous simulation studies of the national liver allocation system. Finally, in Section 2.4 we present previous studies considering optimal usage of organs for transplantation.

2.1 MARKOV DECISION PROCESSES

This section introduces the methodology that is used to formulate the problem. The notation in this section is mostly from Puterman [125]. A collection of objects, (T, S, A_s, $P_t(.|s,a)$, $r_t(s,a)$) is referred to as a discrete-time Markov decision process in which $T = 1, \ldots, N$ are the *decision epochs*, that is, the set of the points of time at which decisions are made; S is the *state space*, that is, the set of observations which the system occupies at each decision epoch and that describe the system; for any state $s \in S$, A_s is the *action space*, that is, the set of possible actions that the decision maker can take at state s; $P_t(.|s,a)$ are the *transition probabilities*, that is, the probabilities that determine the state of the system in the next decision epoch; and $r_t(s,a)$ is the *reward function*, that is, the result of taking action at each state. A *decision rule* is a procedure for action selection from A_s for each state, namely, $d_t(s) \in A_s$. We can drop the index s from this expression and use $d_t \in A$ which represents a decision rule specifying the actions to be taken at all states, where A is the set of all actions.

A *policy* δ is a sequence of the decision rules to be used at each decision epoch and defined as $\delta = (d_1, \ldots, d_{N-1})$. A policy is called *stationary*, if $d_t = d$ for all $t \in T$. A *Markovian* decision rule is one which depends only on the state and action taken at the last decision epoch, whereas a *history-dependent* decision rule depends on the past history of the system represented by the sequence of previous states and actions. A decision rule prescribing exactly one action for each state is called a *deterministic* decision rule whereas a *randomized* rule is one in which there is a probability distribution over the actions that can be taken at each state. There are four types of policies in general: Markovian and deterministic, Markovian and randomized, history-dependent and deterministic, history-dependent and randomized. In this study, we consider only Markovian and deterministic decision rules because it can be shown that when the immediate rewards and transition probabilities depend only on the current state of the system, there exists a corresponding optimal deterministic and Markovian decision rule [125].

MDPs are classified according to the time horizon in which the decisions are made: finite or infinite horizon. They can be further classified into two categories with respect to the existence of a discounting factor for future rewards: discounted or undiscounted. There are three main performance or optimality criteria used for evaluating different decision rules or policies: expected total reward criterion, expected total discounted reward criterion and average reward criterion [125].

The objective of solving an MDP is to find the policy that optimizes the expected reward for a given optimality criteria. For a finite horizon problem, average total reward and expected total reward criteria give the same optimal policies [125]. In the absence of a discounting factor, if we let $u_t^*(s_t)$ be the optimal value of the expected total reward when the state at time t is s and there are $N - t$ periods to the end of time horizon, then the optimal value function and the optimal policy can be computed by iteratively solving the following recursive equations which are also called *Bellman equations*:

$$u_N^*(s_N) = r_N(s_N) \text{ for all } s_N \in S,$$

and

$$u_t^*(s_t) = \sup_{a \in A_s} \left\{ r_t(s_t, a) + \sum_{j \in S} P(j|s_t, a) u_{t+1}^*(j) \right\} \text{ for } t = 1, \ldots, N-1,$$

where $r_N(s_N)$ denotes the terminal reward that occurs at the end of the process when the state of the system at time N is s. At each stage, we choose

$$a^*_{s_t,t} \in \arg\sup_{a \in A_s} \left\{ r_t(s_t, a) + \sum_{j \in S} P(j|s_t, a) u^*_{t+1}(j) \right\} \text{ for } t = 1, \ldots, N-1,$$

where $a^*_{s_t,t}$ is a decision rule maximizing the expected total reward at time t for state s [125]. The most common method used for solving finite-horizon problems, the *backward induction* method uses the above equations to obtain the optimal policies.

The most commonly used optimality criterion for infinite-horizon problems is the expected total discounted reward. The expected total discounted reward, or value, of a policy δ is defined as

$$V^\delta_\lambda(s) = \lim_{N \to \infty} E^\delta_s \sum_{t=1}^{N} \lambda^{t-1} r(X_t, Y_t),$$

where $0 \leq \lambda < 1$ is a discounting factor, X_t represents the decision rule applied to state s and Y_t represents the system state at time t [125].

In a general MDP model, the following assumptions guarantee the existence of optimal stationary policies: stationary rewards and transition probabilities, bounded rewards, discounting with λ where $0 \leq \lambda < 1$ and discrete state and action spaces. Under these assumptions, it can be shown that the value of any stationary policy d^∞ can be computed by solving the following set of equations [125]:

$$V^{d^\infty}_\lambda(s) = \left\{ r_d(s) + \sum_{j \in S} \lambda P_d(j|s) V^{d^\infty}_\lambda(j) \right\} \text{ for } s \in S,$$

where $V^{d^\infty}_\lambda(s)$ is the value of policy d^∞ at state s, $r_d(s)$ is the reward obtained under policy d^∞, and $P_d(j|s)$ is the probability that the system will move to state j at next decision epoch given that the system is in state s and decision rule d^∞ is applied. It can also be shown that there is a unique solution to the following optimality equations giving the optimal policy [125]:

$$V(s) = \sup_{a \in A_s} \left\{ r(s, a) + \sum_{j \in S} \lambda P(j|s, a) V(j) \right\} \text{ for } s \in S,$$

where $V(s)$ is the optimal value of the MDP at state s. The policy maximizing this set of equations is the optimal policy.

MDPs provide powerful models for sequential decision making under uncertainty. The earliest sequential decision making problem in the literature dates back to 1875 when Cayley defined the well known secretary problem [25]. However, the study by Cayley did not resurface until the 1950s when books by Bellman [13] and Howard [80] popularized sequential decision making and MDPs.

Beginning from the 1960s there has been an enormous amount of research in the area of MDPs. MDPs have been applied to a broad class of real-world problems such as inventory control [27, 30, 48, 75, 86, 92, 110], production control [49, 55, 65, 67, 133, 137, 155], finance [54], machine maintenance [134] and medical decision making [1, 74, 81, 87, 98, 102]. Interested readers should refer to books by Puterman [125], Bertsekas [14], White [161], Heyman and Sobel [76], Howard [80] and Bellman [13] for more references and applications.

The most common methods for solving infinite horizon MDPs are value iteration [17, 125, 135], policy iteration [13, 80] and linear programming [43]. However, these methods fail to work well as the MDP problem becomes large, which is often referred as the "curse of dimensionality" [13, 33]. Therefore, many researchers consider efficient solution techniques for solving MDPs. These include variants of the basic methods such as modified policy iteration [113], relative value iteration [160], and approximation techniques such as *neuro-dynamic programming*. Neuro-dynamic programming techniques allow systems to learn how to make good decisions by observing their own behavior and to use their built-in mechanisms for improving their actions through a reinforcement mechanism [15]. In the artificial intelligence community, the name *reinforcement learning* is also used to refer to neuro-dynamic programming [15].

Researchers who formulate MDPs also seek to establish the existence of structured optimal policies. Such policies are appealing to the decision makers and are easier to implement. Furthermore, problems that are known to have special structures may be solved faster by using specialized algorithms. Examples of structured policies are (s, S) policies in inventory models [122, 168] and control-limit policies in queueing control models [12, 125]. A control-limit policy consists of decision rules of the form

$$d_t^*(s) = \begin{cases} a_1 & s < s^* \\ a_2 & s \geq s^*, \end{cases}$$

where a_1 and a_2 are distinct actions and s^* is a control limit [125]. This policy can be interpreted as follows: if the state of the system is less than s^*, it is optimal to take action a_1, otherwise it is optimal to take action a_2. In many areas of application, such as maintenance optimization [11, 29, 120, 150], inventory theory [71, 158], queueing [20, 42, 157, 159], and control theory [96], authors derive sufficient conditions to ensure the existence of an optimal control-limit policy.

In MDP models, we assume the state that the system occupies at each decision epoch is completely observable. However, in some real-world problems, the actual system state is not entirely known by the decision maker, rendering the states only partially observable. Such MDP models are known as *partially observable Markov decision process (POMDP)* models. In addition to the basic components of a MDP model, a POMDP model contains (i) an *observation set* which the decision maker observes at each decision epoch rather than an actual state space, and (ii) a probability function of this space, the state space and the action space. In other words, the decision maker can make inferences about the actual system state using the partial observation that he/she has.

POMDPs were first introduced in the 1960s [6, 7, 51]. Later work [138, 139, 140] addressed the computational difficulties associated with POMDPs and developed algorithms specific to POMDPs.

A finite action/state space POMDP model can be analyzed by transforming it to an equivalent continuous-state Markov decision process in which the system state is a probability distribution on the unobserved states in the POMDP, and transition probabilities are derived through Bayes' rule [125]. Therefore, standard solution techniques used for solving MDPs can also be used for solving POMDPs. In addition to those, there are solution algorithms specific for PODMPs such as Sondik's one-pass and two-pass algorithms [139], Cheng's relaxed region and linear support algorithms [31] and Cassandra's witness algorithm [24]. More information about exact and approximate algorithms for solving POMDPs is available in Monahan [112], as well as in Cassandra [24].

2.2 MDP APPLICATIONS IN MEDICAL DECISION MAKING

This section summarizes the previous applications of MDPs to medical decision making problems. MDPs have been widely used to solve production control, finance, logistics, machine replacement, and inventory control problems but are not very common in medical decision making. Although sequential decision making models are frequently needed in medical decision making problems, using decision trees to solve the arising decision problems is more popular than MDP modeling. However, decision trees become very inefficient as the number of recurrent events increases. Moreover, decision trees cannot be used to formulate problems that have embedded decision nodes [129]. Recent research applies MDPs to medical decision making problems successfully. In this section, we present some of these studies. Most of the material in this section comes from Schaefer et al. [132].

Lefèvre [98] uses a continuous-time MDP formulation to model the problem of controlling an epidemic in a closed population of N people. In his model, the state of the system is described by the number of people infected at time t. Possible actions include quarantine and medical care programs that could change the number of infected people by directly influencing the birth and death rates. Each action has a certain cost and the objective in his decision model is to minimize the total expected discounted cost in the infinite horizon.

Hu et al. [81] address the problem of choosing a drug infusion plan to administer to a patient using a POMDP model. The internal states in this problem are patient parameters that cannot be observed directly, such as the anesthesia concentration in the blood and the clearance rate of the drug. They use approximation techniques to obtain suboptimal though useful treatment options. They define six different policies for adaptive control of drug concentrations in compartment models and compare their performances using a simulation model.

Magni et al. [102] present an MDP model for optimal timing of intervention for treating hereditary spherocytosis disease which destroys red blood cells. The state of the patient is described through the state of the gallbladder and the state of the spleen. The authors consider gallstone natural history, risk of surgical mortality, and natural causes of death in deriving transition probabilities. They use a utility function based on the quality-adjusted

life years (QALYs) for representing rewards, and the objective is to maximize QALYs by deciding the time of the intervention.

Hauskrecht and Fraser [74] use a POMDP framework to model and solve the problem of treating patients with ischemic heart disease (IHD). IHD results from an imbalance between the supply and demand of oxygen to the heart. The state of the patient is described by such variables as the level of coronary heart disease, ischemia level, stress test result, history of coronary artery bypass surgery and angioplasty and so on. Possible actions include medication, surgery (such as angioplasty or bypass surgery), investigative procedures (such as an angiogram or stress test), or doing nothing at all. The internal states of the PODMP model consist of the level of coronary artery disease, i.e. the ischemia level, because this is not directly observable. On the other hand, variables such as the level of chest pain are directly observable. Hauskrecht and Fraser develop an infinite-horizon discounted POMDP model that seeks a treatment strategy that minimizes total lifetime costs, and they use heuristics to solve this problem. Overall, their POMDP formulation is very effective and efficient in generating good treatment strategies for IHD [132].

Ivy [87] considers the problem associated with detection and treatment of breast cancer. She formulates the problem using a POMDP model which determines when to recommend a mammogram and the form of optimal treatment. Alternative treatments include lumpectomy and mastectomy. The goal is to minimize the total expected cost over a patient's lifetime, where costs are based on the patient's condition, tests, and treatment options. She solves this problem using real data.

Ahn and Hornberger [1] provide a Markov decision process model that considers the accept/reject decision of a patient when there is a kidney offer. We describe their model in detail in Section 2.4.

While there have been very few MDP applications to medical decision making, such recent successful applications suggest that MDPs might provide powerful tools for clinical decision making and will be more popular in the near future.

2.3 PREVIOUS SIMULATION MODELS ON LIVER ALLOCATION SYSTEM

This section describes previous simulation models of the liver allocation system and motivates the need for a natural history model of end-stage liver disease. Several discrete-event simulation models have been developed to estimate the effect of various allocation policies. UNOS and the Pritsker Corporation use the UNOS Liver Allocation Model (ULAM) to address issues related to the liver allocation system [124]. The model considers the patient listing process, the organ availability, and the UNOS matching criteria. Post-transplant survival is estimated using Cox proportional hazards models [38]. The primary goal of the ULAM model is to consider the effects of moving from a regional to a national waiting list. Their model considers the old liver allocation model in which there are four statuses assigned to patients for determining their medical urgency.

One limitation of the ULAM model is that a patient's medical history prior to transplantation is given by the rate at which patients transition between predefined priority statuses. The model is therefore unable to consider more general measures of health status. Considering the multiple changes in the liver allocation system recently, flexibility is a significant factor in any simulation model of the liver allocation system. For instance, the medical statuses of the patients that they use for transitions are no longer used by UNOS, and therefore the ULAM model may no longer apply.

It is also important to consider that since a patient can be listed for several years, her health often changes during the time between listing and transplantation. In the ULAM model, patient health can only change by moving from one of the priority states to another. For example, in these models a patient can move from one status to another, but the models keep no specific physiological data other than this.

The CONSAD Corporation [37] uses a simulation model to ascertain the effects of a national distribution scheme. Since this model also measures the natural history of the liver disease by rate of progression through existing priority levels, the model's ability to predict changes in patient health depends on the success of the specific allocation scheme. Therefore,

although this model makes similar assumptions and uses almost identical data as the ULAM model, it reaches slightly different conclusions.

Roberts et al. [95, 128, 136] develop a simulation model to overcome the issues that arises in the other simulation models described above. They design a biologically based discrete event simulation to test proposed changes in allocation policies. Their simulation model consists of five core modules: the patient generator, the organ generator, the pre-transplant natural history, the matching algorithm, and the post-transplant survival. Section 4.2 describes the pre-transplant natural history module in detail. They use mathematical models of disease progression and post-transplant survival to track patients' health at the individual level. Unlike the other simulation models, they model the changes in the medical status of patients independently from the priority statuses, which makes the model sufficiently flexible enough to evaluate a user-defined liver allocation system. Their model is clinically based and modular, so that they can incorporate possible changes to natural history, post-transplant survival, patient prioritization, and the allocation policy.

2.4 PREVIOUS RESEARCH ON THE OPTIMAL USE OF ORGANS

This section describes the studies on the optimal use of cadaveric organs for transplantation. David and Yechiali [40] give perhaps the first model for a problem similar to the problem under consideration. They consider when a patient should accept or reject an organ for transplant. They first assume that organs arrive at fixed time intervals and provide a time-dependent control-limit optimal policy. Next, they consider the case where the arrival of organs is a renewal process and patient health is always deteriorating. Under these assumptions, they show that the optimal control limit is a nonincreasing continuous function, so that a patient becomes more willing to accept lower quality organs as time progresses. Unfortunately, they make many assumptions that do not reflect the actual organ allocation system. For instance, in one of their models the organ arrival rate decreases as time passes. In the actual configuration, a patient should see the organ arrival rate increase as she wins more tie-breakers for length of time on the waiting list. Furthermore, patient health does not,

in fact, monotonically deteriorate [3]. They also do not consider the waiting list, nor do they consider the actual matching criteria.

In a later work, David and Yechiali [41] consider allocating multiple organs to multiple patients. They consider the case where there are more organs than patients, the case where there are equal numbers of each, and the case where there are more patients than organs. The health of the patient is constant, and a rejected organ is never offered to another patient. It is not clear that any of the data that are used are based on clinical observation, and the authors do not provide any computational results from the proposed algorithms. Moreover, they do not consider changes in the patient health. Most importantly, there are serious obstacles to implementing the optimal policies they obtain because they ignore the existence of the waiting list; therefore, they do not address the issue of who is given the priority for an organ offer among patients on the waiting list.

Righter [126] considers a resource allocation problem in which there are a finite number of activities each of which requires a resource, where resources arrive according to a Poisson process. Her model can be applied to the kidney allocation problem, where resources represent the organs and activities represent the patients. She assumes that the reward of assigning a resource to an activity is the product of the resource value and the activity value. She presents several structural properties of the optimal value function such as monotonicity. She does not provide any computational results. She does not consider the effect of the physiology of the patient on the transplant success and post-transplant survival rates [127, 151]. Her model does not consider the effect of the waiting list on the organ arrival rates and the organ allocation process.

Hornberger and Ahn [77] describe a model for kidney transplantation in which the goal is to decide eligibility of the patient for transplantation of a donated kidney. They compute the one-year survival rate of the patient when she accepts the offered organ and receives the transplantation, and this survival probability is used to determine the eligibility of the patient for transplantation with that organ. They consider the effect of the waiting list on the probability of receiving a future organ, and estimate one-year survival rates by performing a Weibull regression on real survival data. They conclude that using minimum eligibility standards improves total quality-adjusted life years for all groups of patients considered.

They do not consider that a patient, by remaining on the waiting list, may gain points in tie-breakers. Moreover, they do not clearly state how the waiting list is incorporated into their model. Furthermore, their model does not consider the fact that the patient's health changes and usually deteriorates over time. Finally, they do not consider the physiology of the patient, which certainly has effects on the transplant success and post-transplant survival rates [127, 151].

Ahn and Hornberger [1] provide an MDP model for deciding which kidneys would maximize a patient's duration and quality of life. In their model, the patient is involved in the process of determining a threshold kidney quality value for transplantation. A patient may experience one of five states: alive on dialysis and waiting for transplantation; not eligible for transplantation; received a functioning renal transplant; failed transplant; and death. Due to the small number of states, they provide an exact analytical solution for threshold kidney quality. Due to insufficient data, the authors use months as their decision epochs; however, it is frequently the case that patients make these decisions more often. In both papers, patient health is modeled very simply and does not include the physiological state of the patient. However, as Valenta [151] finds, the physiological state of a liver patient has effects on the post-transplant survival rates of the patients, and therefore, their framework might not give accurate results for the liver patients.

Zenios et al. [166] develop a Monte Carlo simulation to compare various kidney allocation policies. They compare a fundamental trade-off in organ allocation: efficiency and equity. Patient survival rates and quality-adjusted life expectancies are used as performance criteria for efficiency. Equity is measured by average waiting time to transplant for patients in different ethnic groups. They suggest alternative allocation schemes and show that several of these schemes outperform the current UNOS allocation plan on several criteria. In this simulation model, a patient must accept a liver allocated to her, and patient welfare is only considered in the aggregate.

Zenios et al. [165] consider the problem of finding a best kidney allocation policy with three criteria of maximizing total quality-adjusted life years (QALYs) and minimizing two measures of inequity. The first measures equity across various groups in terms of access to kidneys and the second measures equity in waiting times. They model this problem as a

continuous-time, continuous-space deterministic fluid model which does not have an obvious closed-form solution. They provide a heuristic dynamic index policy for the problem. As with [166], they only consider aggregate measures of patient welfare and do not permit patients to reject kidneys. In both papers they do not consider the dynamics of patient health.

Su and Zenios [142] consider the problem of allocating kidneys to the transplant candidates who have the right to refuse the organs. They use a sequential stochastic assignment model to solve different variants of this problem. First, they consider the case when the patient does not have the right to reject an organ offer and the number of the organs offered is the same as the number of patients on the waiting list. They then consider the problem of allocating kidneys to the patients when the patients have the right to refuse an organ offer and measure the effects of this patient autonomy on the overall organ acceptance and rejection rates. In both problems, the objective is to maximize total reward, where reward is a function of patient health and organ type. For both problems, they derive structural properties of the optimal policies under different reward functions; however, their model has several shortcomings. First, they make unrealistic assumptions such as the number of organ offers is equal to the number of patients on the waiting list. As Table 1.1 shows, there is an imbalance between the number of organs and the number of patients on the waiting list [149]. Second, they do not consider the effect of the waiting list on the number and quality of the organs that are offered to the patient. Furthermore, they do not consider the dynamic nature of patient health and its effect on post-transplant survival and accept/reject decisions. For example, their model does not consider mortality risks for the patients. Third, the performance measures they use might not give the true optimal solution for the patients. For example, they use one-year and five-year graft survival rates and number of wasted organs as their performance measures, which is equivalent to assuming that the patient is indifferent to the timing of organ transplantation and she always benefits from having the transplant as soon as she gets the organ offer. However, as we demonstrate in Chapter 5, waiting longer might give the optimal result for the patient, even if she receives an excellent organ offer. Therefore, using total expected life days or quality-adjusted life days to measure the overall effectiveness of an organ allocation system should be considered for a complete model. As

a result, as they acknowledge, their model fails to capture many important aspects of the kidney allocation problem [142].

Howard [79] presents a decision model in which a surgeon decides to accept or reject a cadaveric organ based on the patient's health. He does not provide any solutions to this decision model. Instead, he provides statistical evidence that explains why a transplant surgeon may reject a cadaveric liver offer. He also discusses the trends in organ procurement in light of his findings.

This dissertation differs significantly from the previous studies. First, this research addresses livers, rather than kidneys. ESLD are different from the end-stage renal diseases (ESRD) because the only viable therapy for ESLD patients is organ transplantation whereas there are alternative therapies such as dialysis for ESRD patients. Second, this research addresses the decision faced by a patient awaiting an organ transplant, which differs from the work of Zenios et al. [165, 166] and Su and Zenios [142] in that no global allocation scheme is proposed. Such research must consider many factors, including ethical issues such as equity among various minority groups [26] and prioritizing liver patients for transplantation [26] and political issues such as states' rights [146]. Because this research views and solves the problem from the patient's perspective, it does not need to consider such issues. Third, unlike previous researchers [1, 40, 41, 77, 126, 142, 165, 166], this research uses detailed models of patient health utilizing real data gathered from clinical observation. Fourth, unlike David and Yechiali [40, 41], this model is not limited by unrealistic assumptions such as the number of organs is greater than or equal to the number of transplant candidates. Finally, this research considers the effects of the waiting list on the organ arrival rates and the decisions of the patients.

As a result, this study describes a model that can be implemented immediately. The implementation of the proposed models does not require unrealistic assumptions about data, nor does it propose or require a change in the current allocation scheme. The next chapter presents the MDP formulation of the decision model.

3.0 PROBLEM FORMULATION

The state space consists of two parts: the liver offered, if any, and the patient health. The state space component corresponding to the liver offered includes all possible liver types that can be offered to the patient, including a state corresponding to a no-offer state. The state space component corresponding to patient health consists of all transient health states, and absorbing states corresponding to pre-transplant death and post-transplant states. We assume that the quality of the organ and the health of the patient may be discretized.

The decisions are made when patient health changes or a liver is offered to the decision maker, which could occur at any moment in time, requiring a continuous time model. However, we approximate the continuous time process by using a discrete-time Markov decision process [125].

There are three possible actions that the patient can take at any decision epoch. She can choose the "Transplant the cadaveric organ" option if one has been offered to her; she can choose the "Transplant the living-donor organ" option if such an organ exists; or she can decline the liver offer(s) and choose to "Wait" for one more time period. When the patient chooses either of the "Transplant" options, she quits the process and moves to the post-transplant state.

There is a consensus that in any sequential decision making context, it is necessary to state all future effects in terms of their present value to the decision maker [152]. Weinstein and Stason [156] point out that the future life years should be discounted because they are being valued relative to dollars and since a dollar in the future is discounted, so must a life year. Keeler and Cretin [91] present a cost-effectiveness study that shows that failure to discount future benefits such as improvements in future life years implies that health-care programs are always improved by delay. On the other hand, the selection of the appropri-

ate discount factor is controversial [22, 39, 62, 70, 78, 90, 94, 108]. For instance, Harvey [72] argues that both psychological evidence and common political observation imply that constant-rate discounting is unrealistic and proposes the usage of nonconstant discount rates. On the other hand, Gold et al. [68] point out some theoretical problems with using nonconstant discount rates and suggest using the same discount rate that is used to express future dollars in present value. More information about selection of the appropriate discount rate to the future health benefits is available in [28, 61, 68, 152].

In this study, we apply a stationary discount rate to future rewards. We do not lose any generality by applying a discount rate to the future rewards. Because there is no reward associated with remaining in the absorbing states, it can be shown that there exists a stationary optimal policy for the discounted [13, 125], as well as the undiscounted case [125].

We assume that the transition probabilities and rewards are stationary. The following notation is common to all three models.

Notation

- $T = \{1, ...\infty\}$: time periods.
- λ: the discount factor, $0 \leq \lambda \leq 1$.
- h_t: patient's health level at time $t \in T$. We assume that there exists a complete ordering of the health states. We use the MELD scores to represent patient health in our computational experiments, which satisfies this assumption.
- ℓ_t: quality of the liver offered to the patient at time $t \in T$. We assume that there exists a complete ordering of the liver offer states. The liver state definition employed in our computational experiments satisfies this assumption as a result of our method for classification of the livers, which is described in Section 4.4.
- $s_t = (h_t, \ell_t)$: the state of the process at time $t \in T$.
- S_H: health state space, i.e. $S_H = \{1, \ldots, H+1\}$, where $H+1$ represents death.
- S_L: organ state space, i.e. $S_L = \{1, \ldots, L+1\}$, where $L+1$ represents the case that no liver is offered.
- S: state space, i.e. $S = S_H \otimes S_L$.
- ℓ_{LD}: living-donor liver quality.

- $a^*(s)$: optimal decision when the state is s, which represents accepting the cadaveric offer or equivalently quitting the process if it is '$\mathrm{T_C}$', accepting the living-donor liver or equivalently quitting the process if it is '$\mathrm{T_{LD}}$' and waiting for one more period or equivalently continuing the process if it equals 'W'.
- $Q(h, \ell, T_C)$ = total expected discounted reward of the patient when the patient has health level h and a liver described by ℓ is transplanted.
- $r(h, \ell, T_C)$: total expected discounted reward if the patient accepts the liver described by ℓ while in health state h.

$$r(h, \ell, T_C) = \begin{cases} 0, & \text{if } h = H+1 \text{ or } \ell = L+1, \\ Q(h, \ell, T_C), & \text{otherwise.} \end{cases}$$

Note that $r(h, \ell, T_C)$ is also a function of the liver quality and patient type, i.e. gender and blood type. However, since we assume these factors are fixed, we suppress this dependency for notational convenience.

- $r(h, \ell_{LD}, T_{LD})$: total expected discounted post-transplant reward that the patient accrues when she is transplanted with the living donor and her health is h at the time of the transplantation. Note that $r(H+1, \ell_{LD}, T_{LD}) = 0$.
- $\rho(h)$: the penalty (disutility) associated with using the living donor when the patient is in health state h at the time of transplantation. We assume that the penalty function depends on patient health, because the utility function of a very sick patient might be different than that of a very healthy patient.
- $r'(h, \ell_{LD}, T_{LD})$: net total expected discounted post-transplant reward that the patient accrues when she is transplanted with the living donor and her health is h at the time of the transplantation, i.e. $r'(h, \ell_{LD}, T_{LD}) \equiv r(h, \ell_{LD}, T_{LD}) - \rho(h)$.
- $r(h, W)$: expected intermediate reward accrued in the current time period when patient health is h and she chooses to wait. Note that $r(H+1, W) = 0$.
- $\mathcal{H}(h'|h)$: probability that the patient will be in state h' at time t+1 given that she is in health state h at time t and the liver is not transplanted at time t. $\mathcal{H}(h'|H+1) = 0, h' \in \{1, \ldots, H\}$ and $\mathcal{H}(H+1|H+1) = 1$.
- $\mathcal{H}$: health transition probability matrix $\mathcal{H} = [\mathcal{H}(h'|h)]$, $h, h' \in S_H$.

- $\mathcal{L}(\ell|h)$: probability that the patient will receive a liver offer ℓ at time t given that she is in state h at time t.
- $\mathcal{L}$: organ arrival transition probability matrix $\mathcal{L} = [\mathcal{L}(\ell|h)]$, $h \in S_H$ and $\ell \in S_L$.
- $\mathcal{P}$: transition probability matrix $\mathcal{P} = [\mathcal{P}(s'|s)]$, $s \in S$, where $\mathcal{P}(s' = (h', \ell')|s = (h, \ell)) = \mathcal{H}(h'|h) \cdot \mathcal{L}(\ell|h')$, $h, h' \in S_H$ and $\ell, \ell' \in S_L$.
- $V(h, \ell)$: maximum total expected discounted reward that the patient can attain when her current health is h and the current liver offered is ℓ.
- $V^i(s)$: the value associated with state s at the ith iteration of the value iteration algorithm.

The above definitions imply that the probability of receiving a liver of type ℓ at time $t + 1$ depends only on the health state at time t and is independent of the type of liver offered at time t. Note also that $r(h, \ell, T_C)$ and $r(h, \ell_{LD}, T_{LD})$ account for the possibility of death during the transplant operation. Furthermore, because patients often need to be retransplanted due to a number of severe post-transplant complications [16, 50, 163], we incorporate the risk and the reward of retransplantation into the post-transplant reward functions.

Figure 3.1 shows the state-transition diagram of the MDP. The optimal solution to this problem can be obtained by solving the following set of recursive equations [125]:

$$V(h, \ell) = \max \left\{ r'(h, \ell_{LD}, T_{LD}), r(h, \ell, T_C), r(h, W) + \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h', \ell'|h, \ell) V(h', \ell') \right\}, \quad h = 1, \ldots, H, \ell = 1, \ldots, L + 1, \tag{3.1}$$

where $r'(h, \ell_{LD}, T_{LD}) = r(h, \ell_{LD}, T_{LD}) - \rho(h)$.

The next chapter describes the sources of the data that we use in our computational tests and three supplementary models that we use to calculate the transition probabilities and the terminal rewards.

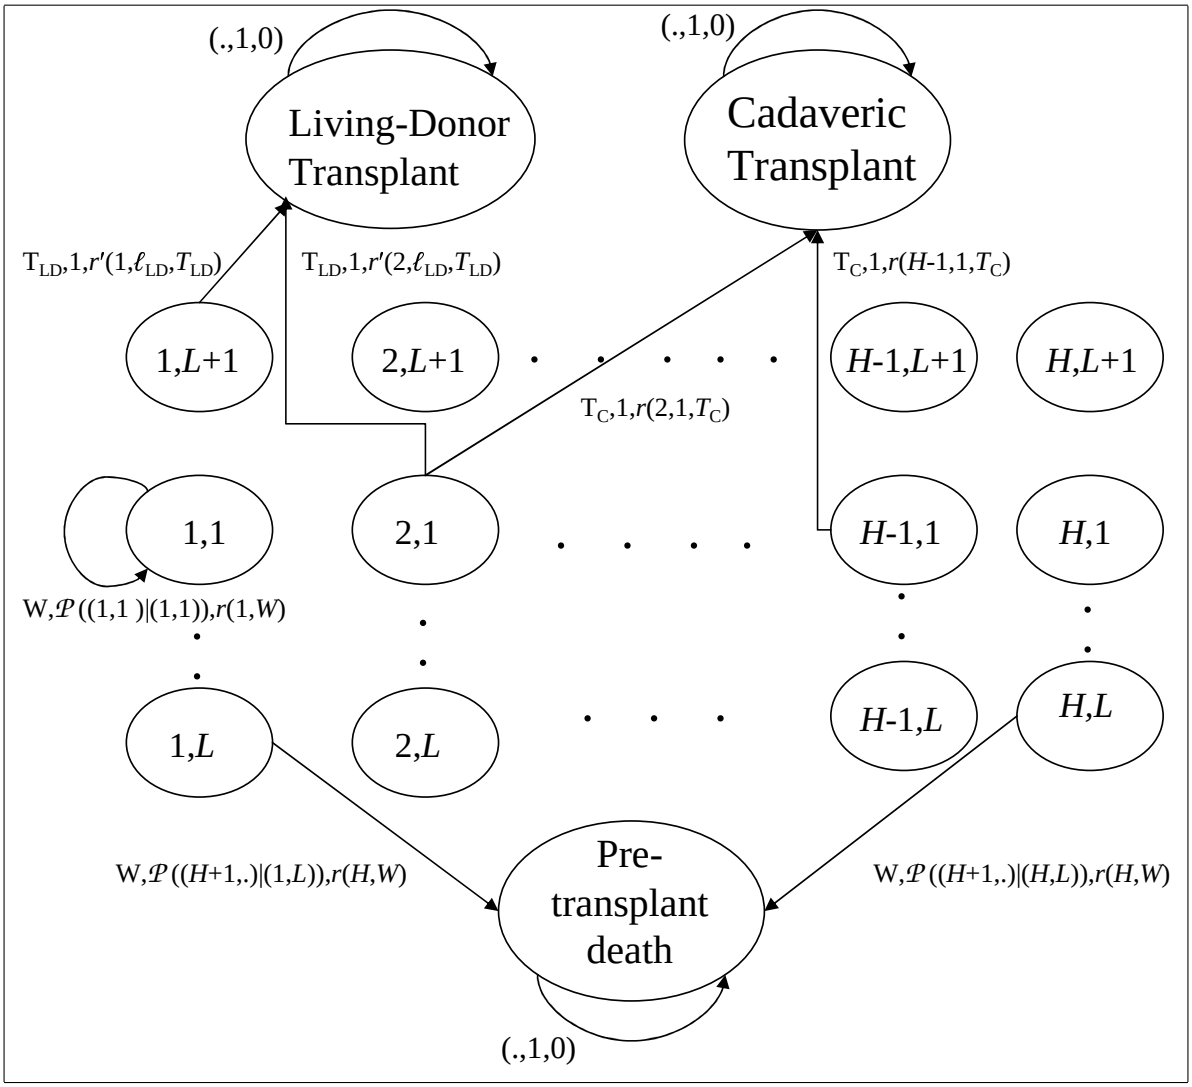

Figure 3.1: State-Transition Diagram of the General Model*

*The labels $a(h,\ell)$,$\mathcal{P}(j|h,\ell)$,$r(h,\ell,a)$ on each arc represent the action taken at state (h,ℓ), the probability that the patient will move to state j when her current state is (h,ℓ) and the reward obtained by taking action a in state (h,ℓ), respectively. Note that this figure does not show all possible transitions.

4.0 DATA SOURCES AND SUPPLEMENTARY MODELS

This research involves estimating model parameters from clinical data. The parameters that must be estimated are $\mathcal{H}$, $r(h, \ell_{LD}, T_{LD})$, $r(h, \ell, T_C)$ and $\mathcal{L}$. In this chapter, we describe the sources of the data and the supplementary models that are used to generate inputs to our decision models. Section 4.1 presents the sources and characteristics of the real clinical data that are used in our computational experiments. In Section 4.2, we describe the methodology that is used to analyze the data and estimate the $\mathcal{H}$. Section 4.3 summarizes the survival model that is used to estimate $r(h, \ell_{LD}, T_{LD})$ and $r(h, \ell, T_C)$. Section 4.4 describes the computation of organ refusal rates and the estimation of $\mathcal{L}$.

4.1 DATA SOURCES

The data come from four sources. The first, *UNOS1*, is a publicly available data set from UNOS that covers 28,717 patients listed for their first liver transplant between 1990 and 1996 and contains data through 1999. UNOS1 is used to determine the post-transplant survival rates. It cannot be used in the development of our model of disease progression because UNOS records included results of laboratory testing done only at the time of transplantation.

The second data set, *UPMC*, comes from The Thomas E. Starzl Transplantation Institute at the University of Pittsburgh Medical Center (UPMC), one of the largest liver transplant centers in the world. Unlike the UNOS data sets, this data set is not publicly available. UPMC is used to estimate the $\mathcal{H}$ matrix as described in Section 4.2 and Chapter 5. We examine the records of 3,009 patients who had ESLD and joined the waiting list between 1991 and 2000. After we exclude patients with incomplete data at listing and patients seen

only once at UPMC, we are left with a sample of 1,997 individual patients. For each patient our data set contains demographic and clinical data, including the type of liver disease that led to ESLD, the results of all laboratory testing done at UPMC, the location of the patient at the time of laboratory testing (at home, in a general hospital ward, or in the intensive care unit [ICU]), information concerning the occurrence of death before transplantation, and information about the existence of clinical covariates such as the presence or absence of encephalopathy.

The third data set, *UNOS2*, was collected between February 27, 2002 and May 31, 2003. We calculate the organ arrival probabilities using UNOS2. The first part of UNOS2 includes information for 25,810 patients waiting for a liver transplant such as region, MELD scores, age, blood type, gender, race, and disease type. The second part of UNOS2 covers information for the cadaveric organ offers such as the patients that each organ has been offered to, the date of the offer, as well as information for the cadaveric donor such as gender, cause of death, age, gender and region number.

The fourth data set, *UNOS3*, was collected between February 27, 2002 and April 30, 2003. We compute the organ refusal rate using UNOS3. This data set summarizes statistics regarding the total number of organ offers and the total number of refusals by recipient MELD score. Four categories for recipient MELD scores are reported: (1) MELD < 11; (2) MELD 11 to 18; (c) MELD 19 to 24; (d) MELD > 25. Unlike the third data set, this data set does not include information about the livers that are not used for transplantation.

4.2 NATURAL HISTORY MODEL

In this section, we describe the methodology that we use to estimate $\mathcal{H}$. Most of this section is from Alagoz et al. [2, 3] and Bryce et al [21]. As described in Section 2.3, previous simulation models of the transplant system that have been used to evaluate allocation policies a priori are inadequate because they fail to estimate disease progression on the basis of clinical or biological factors [37, 124]. Instead, these models estimate disease progression entirely through probability distributions that describe how patients transition from one specific priority status to another.

Estimation of disease progression requires tracking individual patients and updating natural history at any point in time on the basis of biological criteria. A possible method for tracking individual natural histories is to use statistical techniques such as regression to determine how an individual's clinical characteristics change over time. The problem with this approach is that the modeling of individual natural histories from population averages smoothes out the chaotic behavior of specific diseases seen in actual patients. For example, the clinical course of ESLD is often unpredictable and characterized by acute exacerbations and recoveries. Therefore, any simulation model that represents individual patients with ESLD must have the ability to replicate the occasionally chaotic time course of the disease. Methods used to predict the natural history of disease in individual simulation models should balance two important considerations: they should allow for chaotic paths and acute exacerbations in individual patients while still preserving the statistical properties of the natural history in the patient population over time.

This section describes development and validation of our empiric, quantitative natural history model (NHM) of ESLD that estimates the natural history of ESLD and provides quantitative clinical histories of individual patients over time.

4.2.1 The Model

For patients with ESLD, the clinical condition at a given time (t) can be represented by a vector of clinical covariates ($\overrightarrow{X_t}$) that are associated with liver function or transplant success. In our model, this vector contains laboratory values (bilirubin, creatinine, and albumin levels and prothrombin time) and clinical conditions (the presence or absence of ascites, encephalopathy, and mechanical ventilation).

In general, we define a quantitative NHM as a model with the ability to predict the value of a clinical covariate vector at time $t + \Delta t$ given the values of the vector at time t:

$$\overrightarrow{X}_{t+\Delta t} = f(\overrightarrow{X}_t, \Delta t),$$

where $\overrightarrow{X}_t$ is the vector of clinical covariates at time t and $\overrightarrow{X}_{t+\Delta t}$ is the vector of these values Δt in the future.

The function f should satisfy several conditions. First, the natural history of disease in individual simulated patients should replicate the stochastic and sometimes chaotic nature of the disease seen in real patients. Second, the natural history of disease in the simulated cohort should mirror the average changes in laboratory values that are observed in a cohort of real patients over a given period of time. Third, the correlations between clinical characteristics and laboratory values in the model should mirror the correlations observed in real patients. Fourth, the pre-transplant survival rates in the model should be similar to those observed in real patients.

Standard statistical estimation techniques are not well suited to satisfy these requirements. For example, regression techniques summarize clinical experiences as the average change over time and eliminate the stochastic nature of real clinical histories. Although time series analysis can create waxing and waning levels of a single variable over time, this type of analysis does not provide a mechanism for simultaneously predicting the time course of multiple correlated variables. Moreover, it does not offer a mechanism for dealing with the data sets that have an inherent sampling bias, which is caused by the oversampling of laboratory values in patients when they are severely ill and undersampling in patients when they are less ill.

The NHM reduces this sampling bias by interpolating additional laboratory values between the existing (real) laboratory values using cubic splines by the method shown in the top portion of Figure 4.1. A function f is called a cubic spline on $[x_1, x_N]$ if the following criteria are met: the function is defined on $[x_1, x_N]$; the function f is continuous and its first and second derivatives exist, both being continuous on the points (knots) of the cubic spline f; and the knots are such that $x_1 < x_2 < \ldots < x_N$, where f is a polynomial with a degree less than or equal to 3 on each subinterval $(x_j, x_{j+1}), j = 1, \ldots, N-1$ [69].

The spline function is estimated by minimizing the penalized least-squares function [69], which consists of two terms: the penalty function and the sum of squares between the observed and expected data points. The addition of the penalty function ensures that the estimated spline function gives the best compromise between its goodness of fit to the data (as quantified by the sum of squares between the observed and expected data points) and its smoothness. The penalty term also is used to assure that the values of the splines do not

exceed certain known limits: for each laboratory value there are minimum possible values, and for some, maximum plausible values, and the simulation routine should only be able to create values between those limits. This technique is called the smoothing spline method. For more information on cubic spline functions, see [69].

For each patient and each laboratory value in the sample, the NHM creates a cubic spline. In the top panel of Figure 4.1, the solid line forming the curve represents the spline, the solid circles represent the observed (actual) bilirubin levels, and hollow circles represent bilirubin levels that are interpolated according to the fitted spline. We use penalized cubic splines with a variable number of knots, so the spline is not required to go through each actual data point. Because interpolated values are calculated at a standard time interval, they provide a large number of time points to sample from, and they reduce the bias of having more observations at a time when patients are ill. Each spline is then decomposed into triplets of three sequential laboratory values (middle panel of Figure 4.1), and merged with the triplets from the same times of the other laboratory values and clinical characteristics (lower panel of Figure 4.1). An individual cubic spline is created for each type of laboratory value (bilirubin, creatinine, and albumin level and prothrombin time) in each patient. The splines are then sampled at regular intervals to obtain a complete longitudinal history of each patient [69].

The spline routines in our model use an individual patient's first and last observed values for a given laboratory test (i.e., the first and last bilirubin levels) as anchors. At each anchor point, the observed and spline-derived values are identical. Because the pace of disease is generally faster when patients are hospitalized, our model samples the splines at daily intervals for patients in the hospital and in the ICU, and it samples the splines at monthly intervals for patients at home. From the spline of laboratory values for a given parameter (i.e., bilirubin levels), the model selects each set of three sequential values and converts them into a time triplet ($t_1 - t_2 - t_3$ triplet), as shown in the middle portion of Figure 4.1. This process is repeated for the other laboratory parameters (i.e., creatinine and albumin levels and prothrombin time) so the model will have time triplets in the same time sequence for all four laboratory parameters, as shown in the bottom portion of Figure 4.1.

Next, the model stratifies the time triplets by liver disease group and patient location. Table 4.1 shows the number of individual patients and the number of spline-interpolated

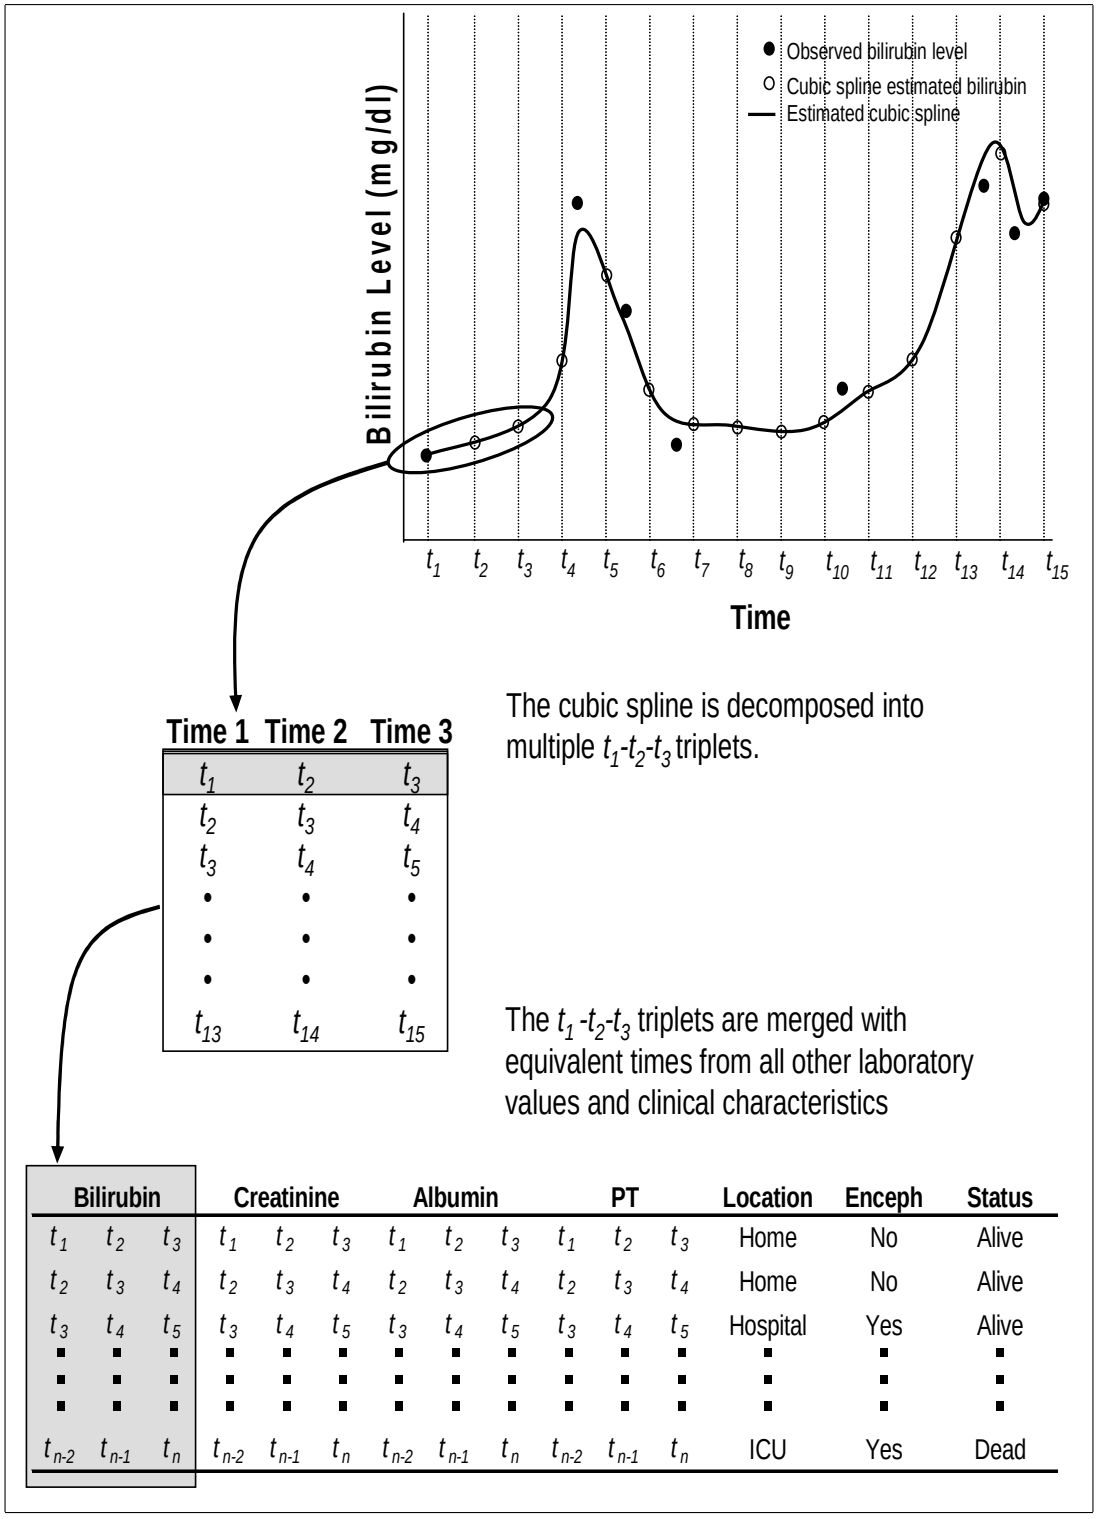

Figure 4.1: Interpolation of Existing Data with Spline Functions [3]*

*PT indicates prothrombin time; Enceph, encephalopathy; and ICU, intensive care unit.

time triplets generated for each disease and location. The data come from the sources that are described in Section 4.1. Sample size prohibits our incorporating all of the types of liver disease in our model, so we aggregate the diseases into five groups: cirrhotic diseases, hepatitis infections, acute liver diseases, cancers, and other liver diseases. These groups were discussed and finalized by the National Clinical Oversight Committee. The diseases included in each group and the members of the Oversight Committee are shown in Table A.1 and Appendix B, respectively. We estimate a disease-specific $\mathcal{H}$ by using these spline-interpolated time triplets that are stratified by liver disease group.

4.2.2 A Simulation Model to Predict the Progression of ESLD

Figure 4.2 describes the method for predicting the progression of disease in an *index patient* with the following laboratory values at the current time (time t): bilirubin level of 8.6 mg/dL; creatinine level of 0.4 mg/dL; albumin level of 1.7 g/dL; and prothrombin time of 24.0 seconds. The values in the previous period (time t-1) determine whether the laboratory values are improving, worsening, or essentially staying the same. To estimate the patient's laboratory values in the next period (time t+1), the model searches the entire list of patients within the same disease group and location to find the first patient who is similar to the index patient in terms of laboratory values and clinical characteristics. The model uses a multivariate nearness measure (defined below) to assess how "near" each patient is to the values of the index patient. When a sufficiently similar patient is found, the model assigns this patient's time t_3 laboratory values to the index patient. These values then become the index patient's new time t values, and the process is repeated the next time the patient's laboratory values need to be updated. Note that in the figure, t_1 represents time t-1, t_2 represents time t and t_3 represents time t+1 for a given observation.

When searching a database for similar patients, the NHM considers a laboratory value to be "near" only if it falls within laboratory-specific upper and lower limits determined by a survey of gastroenterologists and liver transplant intensive care physicians at the University of Pittsburgh. For various levels of each laboratory test, the physicians were asked how much higher or lower the laboratory value would have to be before it would be clinically relevant

Table 4.1: Number of Spline-interpolated Time Triplets [3]

Disease Group and Location	Number of Patients	Number of Interpolated Triplets
Patients with cirrhotic diseases		
At home	869	11,161
In a hospital general ward	834	2,150
In the intensive care unit	493	3,013
Patients with hepatitis infections		
At home	590	7,892
In a hospital general ward	533	8,901
In the intensive care unit	314	1,371
Patients with acute liver diseases		
At home	41	148
In a hospital general ward	69	1,016
In the intensive care unit	59	454
Patients with cancers		
At home	106	1,090
In a hospital general ward	115	2,260
In the intensive care unit	40	294
Patients with other liver diseases		
At home	328	4,467
In a hospital general ward	308	6,147
In the intensive care unit	192	1,297
All disease groups and locations	4,891*	51,661

*Column sums to more than the number of unique patients in the cohort as each patient contributes a record for each location (at home, in general hospital ward, and in ICU).

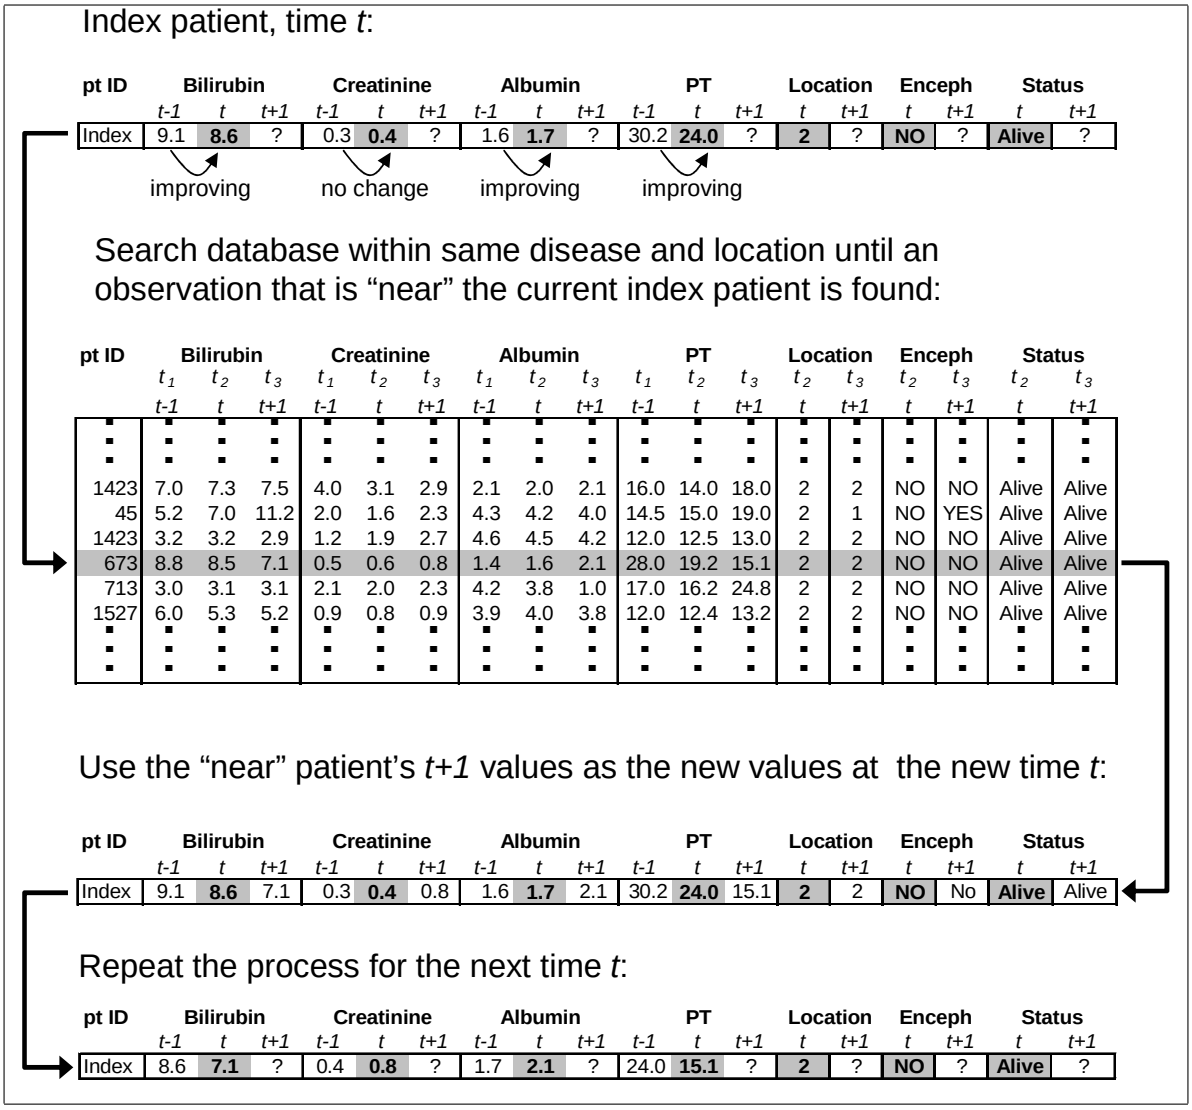

Figure 4.2: Updating a Patient's Health [3]*

*pt ID indicates patient identification number; PT, prothrombin time; and Enceph, encephalopathy.

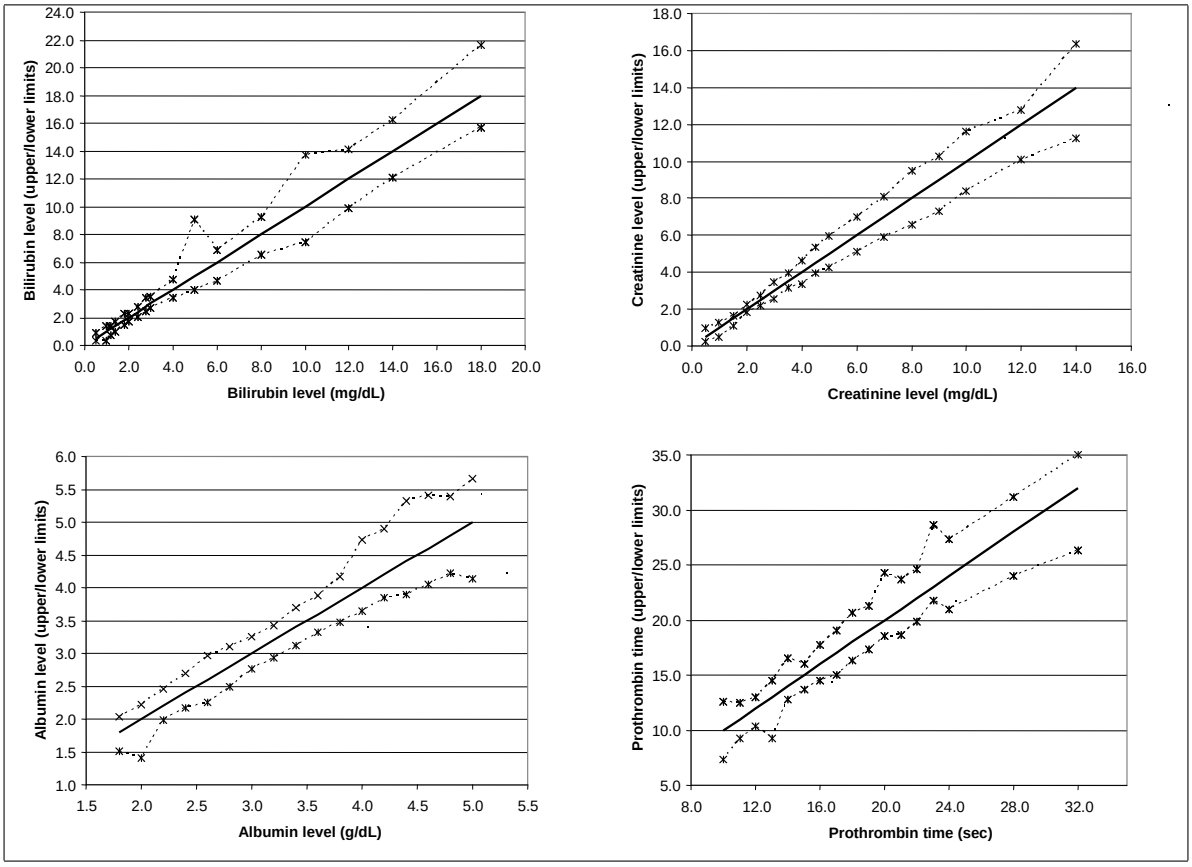

Figure 4.3: "Nearness" Criteria for the Four Laboratory Values Included in the Model [3]

to them in caring for the patient. Because some panel members did not provide values for all levels of the indicated laboratory test, the upper and lower limits are not always monotonic in increasing value. The laboratory values are aggregated and limits average to render the limits smoother and continuous. The results of this survey for the four laboratory values are presented in Figure 4.3. As expected, the limits widen as the level of the laboratory values increases. For example, a given change in the creatinine level is more important when the level is 2.0 mg/dL than when it is 10.0 mg/dL.

The NHM has flexibility in deciding how many of the four laboratory parameters must be "near" to consider the circumstances of two patients to be similar. It also has the option of searching for matches in terms only of the level of laboratory values (a level-based search, or LBS) or searching for matches in terms of both the level and direction of change from the prior time (a direction-based search, or DBS).

We formulate two different DBS strategies. DBS1 seeks a strict match in which all four laboratory values must be both at a similar level and moving in the same direction to be considered close. DBS2 is a less restrictive match in which patients are classified as improving, worsening, or remaining stable, depending on the direction of change in the majority of their laboratory values. In DBS2, a patient is classified as getting worse if three or more laboratory values are getting worse. A patient is classified as being in stable condition if two laboratory values are worsening and the other two are improving. All remaining patients are categorized as getting better. DBS2 assumes that the index patient is moving in the same direction with the "sufficiently close" patient if both patients are classified in the same category.

These search methods can easily be extended to include other quantitative and qualitative clinical criteria for determining the "sufficiently close" patient. If the search fails to find a patient who meets the nearness criteria, the criteria are expanded and the search is repeated.

The NHM has two mechanisms in place to calculate pre-transplant death. The first mechanism is empiric. Because the vital status of each patient at each observation time is recorded, patients who die while waiting for a transplant will have a death indicator in the t_3 position of their last time triplet for location. Thus, if the patient who is found to be "sufficiently close" to the index patient died in the next time period, the index patient

is assumed to be dead at time $t+1$ as well. This approach, the NH death mechanism, will work well only if the actual data set records all of the deaths. Since we used clinical data taken only from our own hospital clinical information system, and because many of our patients receive their primary care at other institutions and are therefore lost to follow-up, our database may not include records of all the deaths, and may therefore underestimate the actual number of deaths.

The second mechanism uses the characteristics of the index patient at each time point to calculate the probability of death during the next time period. Specifically, it uses scores derived from MELD. Using the MELD three-month death probabilities, we first compute the monthly and daily probabilities of death. Then at each time period t, we estimate the probability that the index patient (with his or her particular MELD score) would be dead at time $t+1$. If the patient is not predicted to be dead at time $t+1$, then the NHM looks for a "sufficiently close" patient who is alive and the model proceeds as usual.

4.2.3 Computational Tests

We conduct some computational experiments to validate the NHM. The face validity of the NHM is supported by examination of the model output. Figure 4.4 illustrates the predicted longitudinal histories of two hypothetical patients who initially had identical laboratory values for all four laboratory parameters. Each patient experiences a different history during the course of liver disease. For instance, Patient 1 is stable and out of the hospital for almost a year, is admitted for an exacerbation around day 300, improves and leaves the hospital only to be readmitted for another exacerbation three months later. The patient's liver diseases at this point continue to worsen, and the patient is admitted to the hospital again on day 678 and remains in the hospital until transplant on day 706. In contrast, Patient 2 starts in the hospital, but has a stable course for over three years, when an acute exacerbation and worsening of bilirubin brings the patient back into the hospital. The patient improves, leaves the hospital only to be readmitted on day 1104, and dies during that hospitalization. A detail of the final 16-day hospitalization is shown in the lowest panel: the patient is admitted to the ICU, improved and is moved to the floor, but has acute deterioration between admission

days 10 and 11, is moved back to the ICU, and dies on day 16. Note that these are only example trajectories for liver patients generated by the simulation model; the NHM is able to produce a very large number of individual natural histories.

We simulate patients initially equivalent to patients in the UPMC database using the NHM with all three search techniques (LBS, DBS1, and DBS2) and two different death mechanisms (the NH mechanism and the MELD mechanism), and then we compare the simulation results with the actual UPMC results. Specifically, we compare the following: (1) the average change in laboratory values over time in the simulated cohort versus the UPMC cohort, (2) the correlations between the laboratory values in the simulated cohort versus the UPMC cohort, and (3) the predicted survival rates in the simulated cohort versus the observed survival rates in the UPMC cohort and the national UNOS cohort.

Although we model the results of the disease progression mechanism and the death mechanisms for all five disease groups, here we present the results for only two of these groups: (1) cirrhotic diseases, an example of a chronic liver disease category, which includes primary biliary cirrhosis, primary sclerosing cholangitis, alcoholic liver disease, and autoimmune disorders, and (2) all acute liver diseases.

Tables 4.2 and 4.3 show the average change in laboratory values of the actual UPMC cohort and the simulated cohort for patients with cirrhotic diseases and patients with acute liver disease, respectively. For purposes of comparison, neither the UPMC cohort nor the simulated cohort contain data from patients who had a transplant before the duration of the simulation is exceeded. For example, if a patient had a transplant in month 4, that patient's change in laboratory values would be included in the 3-month cohort, but not the 6-month or 1-year cohort. Therefore, the initial average laboratory values are different for each time horizon. Note also that the change in laboratory values at time t do not include the laboratory values of the patients who died before time t. The table represents only the natural history of those patients who did not die and were not transplanted in the time period. The average change is generally similar in the two cohorts at 3-, 6-, and 12-months; and any differences that are noted between the values in both cohorts are deemed clinically insignificant according to the criteria used to determine if a particular laboratory value is clinically "near" or not. No single search method dominates the others in producing

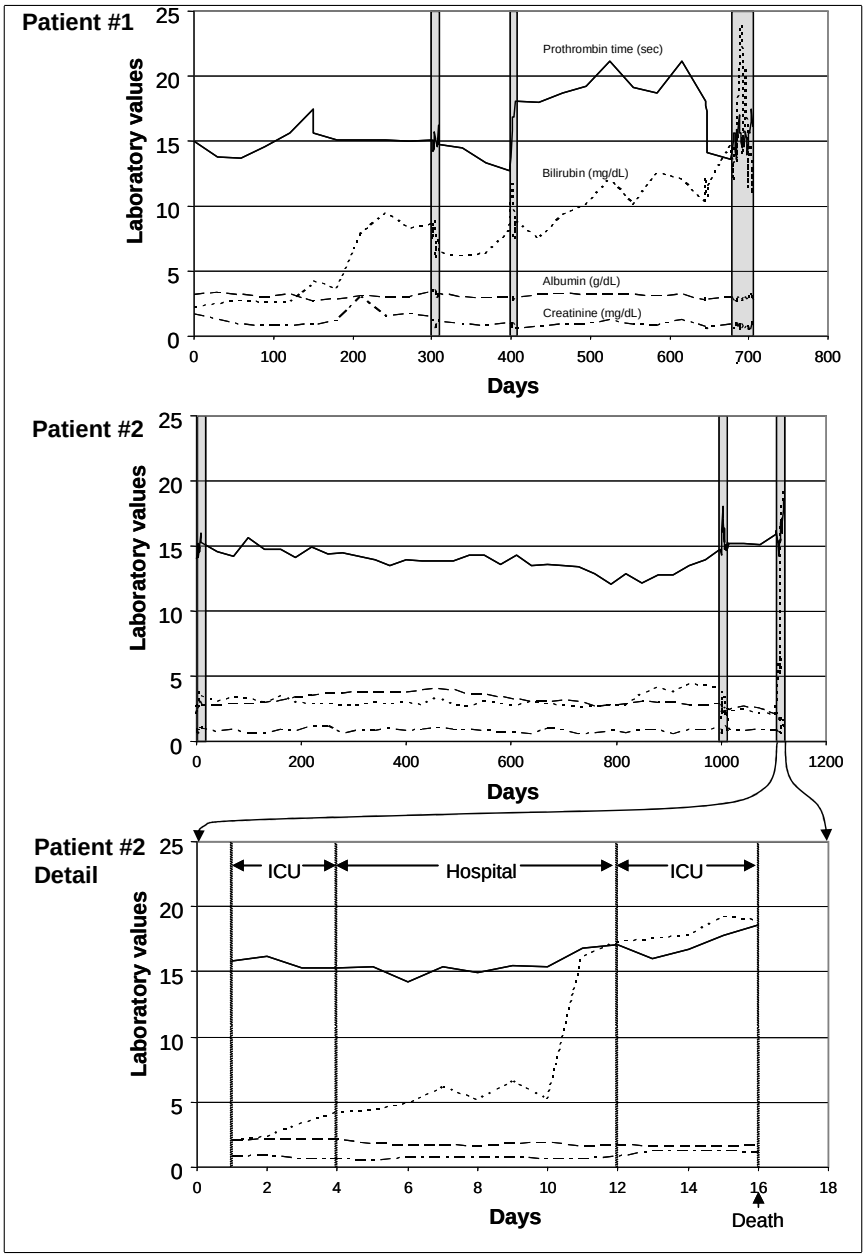

Figure 4.4: Longitudinal Histories of Two Simulated Patients [3]

Table 4.2: Average Change in Laboratory Values of Patients with Cirrhotic Diseases [3]*

		Simulated Cohort		
Disease Group and Laboratory Value	UPMC Cohort	LBS	DBS1	DBS2
Patients with cirrhotic diseases				
Bilirubin level, mg/dL				
3 months	4.2	4.5	2.9	4.8
6 months	2.6	4.4	2.7	4.6
12 months	2.0	4.5	2.4	4.6
Creatinine level, mg/dL				
3 months	1.4	1.0	1.0	1.0
6 months	1.3	1.0	1.0	1.0
12 months	1.3	1.0	0.9	0.9
Albumin level, g/dL				
3 months	3.4	3.2	3.3	3.1
6 months	3.4	3.2	3.3	3.2
12 months	3.3	3.2	3.3	3.2
Prothrombin time, sec				
3 months	14.3	14.2	13.6	14.4
6 months	13.5	14.2	13.5	14.2
12 months	13.3	14.2	13.3	14.2

*UPMC indicates University of Pittsburgh Medical Center; LBS, level-based search; DBS1, restrictive direction-based search; and DBS2, less restrictive direction-based search.

Table 4.3: Average Change in Laboratory Values of Patients with Acute Liver Disease [3]*

		Simulated Cohort		
Disease Group and Laboratory Value	UPMC Cohort	LBS	DBS1	DBS2
Patients with acute liver disease				
Bilirubin level, mg/dL				
3 months	4.3	4.3	1.4	7.2
6 months	3.1	3.9	1.8	6.6
12 months	3.7	3.8	1.9	6.0
Creatinine level, mg/dL				
3 months	1.3	1.2	1.0	1.4
6 months	1.1	1.2	1.0	1.4
12 months	0.8	1.2	0.9	1.5
Albumin level, g/dL				
3 months	3.6	3.2	3.5	3.3
6 months	3.2	3.3	3.4	3.4
12 months	3.2	3.3	3.4	3.4
Prothrombin time, sec				
3 months	14.4	14.6	12.7	14.1
6 months	13.8	14.4	13.1	14.0
12 months	13.3	14.4	13.2	13.9

*UPMC indicates University of Pittsburgh Medical Center; LBS, level-based search; DBS1, restrictive direction-based search; and DBS2, less restrictive direction-based search.

the closest laboratory value changes across disease groups. The performance of the various search methods is not affected by the duration of the simulation. The performance does vary, however, across disease groups, with better results seen in patients with acute liver diseases than in those with cirrhotic diseases. In general, the average actual UPMC laboratory values improve over time because sicker patients with worse laboratory values die or are transplanted and are removed from the cohort.

Table 4.4 shows the correlations between the laboratory values in the UPMC cohort and the simulated cohort of patients with cirrhotic diseases. The simulated natural histories are obtained using LBS and DBS1, DBS2 with the NH death mechanism. Correlations between clinical variables found in the actual data set are maintained in the simulated one.

Figure 4.5 presents 3-, 6-, and 12-month survival rates for the patients in the UPMC cohort with cirrhotic disease used to calibrate the natural history model, and the survival rates of simulated patients using the LBS, DBS1 and DBS2 strategies. The cohort and the simulations contain only patients who are not transplanted during those intervals. We also provide comparison to the pre-transplant survival rates found in the UNOS1 data set, although there is no guarantee that the initial clinical conditions are the same, as initial laboratories at time of listing are not available in the UNOS cohort for that time period. The top panel of Figure 4.5 compares actual survival to simulated survival using the NH death mechanism and the lower panel of Figure 4.5 presents the same results when the simulation uses predicted survival from the MELD score. The simulated results are the average survival at each time point of 100 iterations of simulating the natural history of 602 patients, representing the same size cohort from the UPMC experience. The simulation that closely matches the actual survival when using the NH (empiric) death mechanism and when determining "closeness" with either the LBS or DBS1, DBS2 does not well replicate actual experience. However, when the MELD score-derived death rates are used to predict mortality, all search strategies markedly overestimate pre-transplant death. Regardless of the search technique used, the NH death mechanism produces fewer deaths than does the MELD death mechanism for all of the disease groups (data not shown). As the duration of the simulation increases, the gap grows between the survival rate associated with the NH death mechanism and the survival rate associated with the MELD death mechanism. This

Table 4.4: Correlation of Laboratory Values of Patients with Cirrhotic Diseases [3]*

	Correlation of Laboratory Values			
Cohort and Laboratory Value	Bilirubin Level	Creatinine Level	Albumin Level	Prothrombin Time
UPMC Cohort				
Bilirubin level	1			
Creatinine level	0.329	1		
Albumin level	-0.100	-0.006	1	
Prothrombin time	0.389	0.119	-0.273	1
LBS Simulated Cohort				
Bilirubin level	1			
Creatinine level	0.288	1		
Albumin level	-0.072	0.024	1	
Prothrombin time	0.469	0.198	-0.257	1
DBS1 Simulated Cohort				
Bilirubin level	1			
Creatinine level	0.332	1		
Albumin level	-0.231	-0.024	1	
Prothrombin time	0.494	0.139	-0.431	1
DBS2 Simulated Cohort				
Bilirubin level	1			
Creatinine level	0.409	1		
Albumin level	-0.005	0.161	1	
Prothrombin time	0.432	0.201	-0.140	1

*UPMC indicates University of Pittsburgh Medical Center; LBS, level-based search; DBS1, restrictive direction-based search; and DBS2, less restrictive direction-based search.

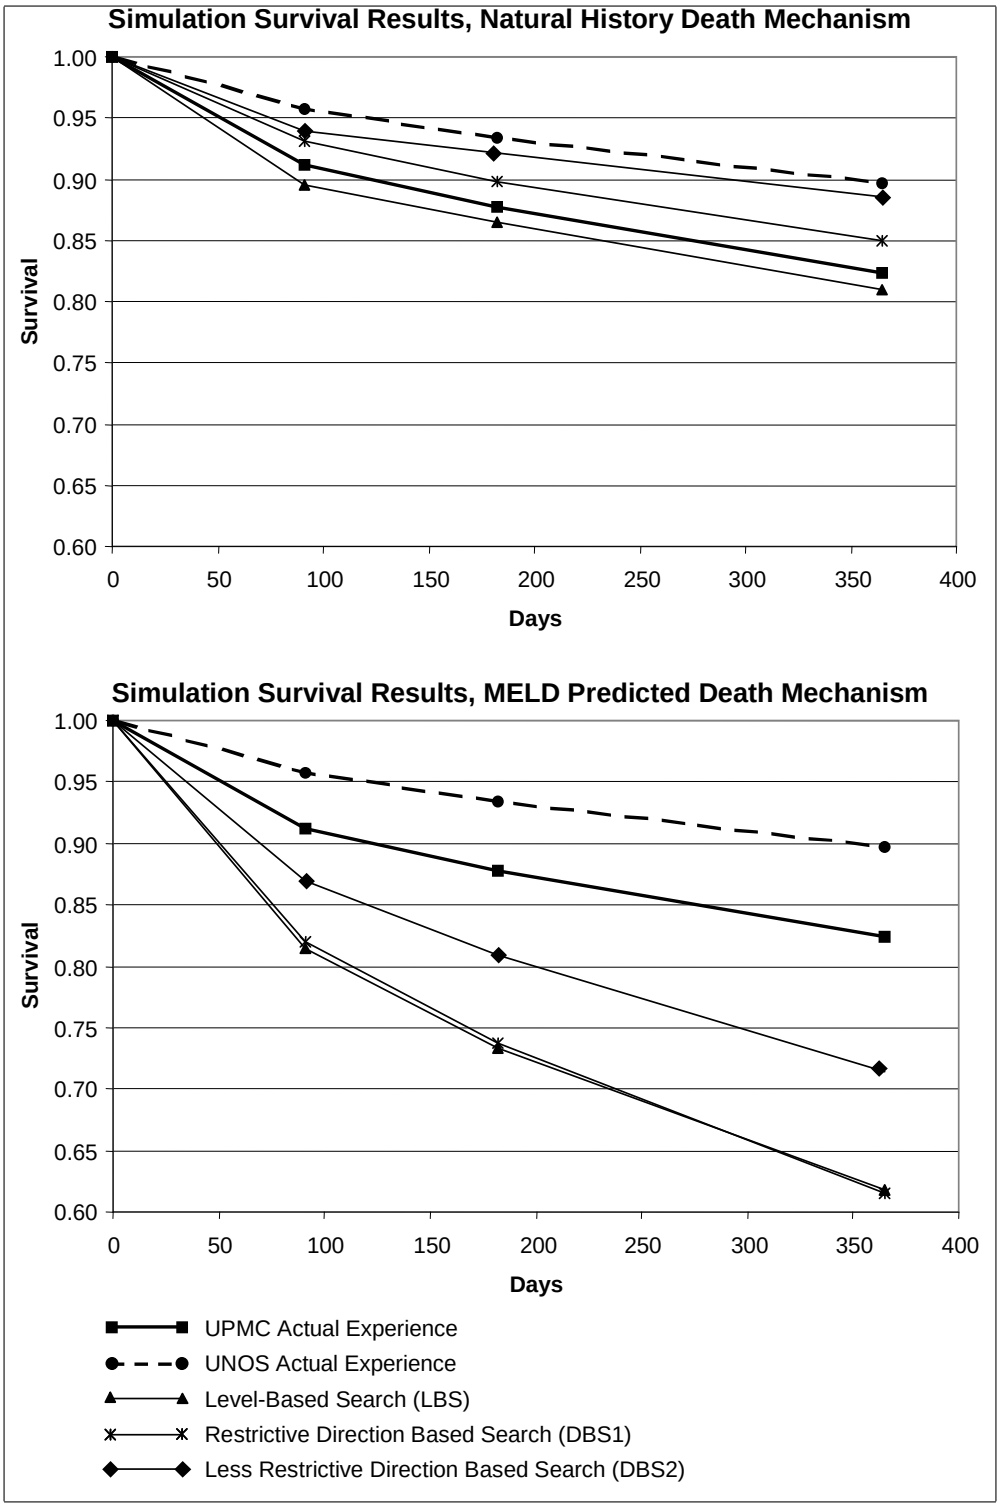

Figure 4.5: Survival Rates in the Simulated and Actual Cohorts [3]

is not surprising and relates to the MELD score itself, which was designed only to predict 3-month mortality [88, 103, 162].

4.2.4 Conclusions and Limitations

The NHM described in this section uses cubic-spline-estimated data from real patients awaiting liver transplantation at UPMC to simulate the laboratory values and clinical status of individual candidates over time, and thereby provides a quantitative natural history of each transplant candidate's disease. Subsequent validity checks show that the NHM is able to simulate the types of acute exacerbations and chaotic disease trajectories that are seen in individual patients, and it is also able to preserve the overall statistical properties of the natural history of ESLD seen in the cohort of real patients over time. Moreover, the NHM is able to predict pre-transplant survival rates that are statistically similar to the rates observed in the UPMC cohort and in the national UNOS cohort. It is interesting that the MELD score, used by UNOS to predict the 3-month risk of death while awaiting transplant, appears to be a poor predictor of long-term survival.

Two factors may affect the generalizability of the NHM. First, because of the intensive data requirements of our model, we are restricted to using data from a single, major university center. Second, it is impossible to directly verify the accuracy of our spline-derived laboratory values. The decision to use data from UPMC and to develop the interval-based spline method both arose from the fact that currently there is no large, multicenter study that collects routine, interval-based data on the natural history of liver disease. UNOS originally collected very little clinical data on candidates at the time of their registration for a transplant, although it began to collect more data after it adopted the MELD criteria for liver allocation in 2002.

Like the current databases for ESLD, the databases for many other diseases are incomplete. We believe that the framework developed in our model for ESLD can be used to generate natural histories for individual patients with other diseases.

4.3 POST-TRANSPLANT SURVIVAL MODEL

Most discussion in this section is from the paper by Roberts et al. [127] and the unpublished dissertation of Valenta [151]. Roberts et al. [127] describes a disease-specific survival model that predicts post-transplant survival rates of an ESLD patient given a set of clinical characteristics. The UNOS classification scheme includes more than sixty codes for the ESLD, many of which include few patients for a disease-specific analysis. Therefore, they use the disease classification described in Table A.1.

They first ask the National Oversight Committee to determine the set of candidate covariates that might affect the survival. They then test each candidate variable in single variable analyses, using a log-rank test or Cox proportional hazards model. They select those variables whose association with outcome had a significance level of 0.1 or smaller. They use these variables to construct multivariable survival models. Finally, they construct disease-specific survival estimates using a Cox proportional hazards regression model [127, 151]. We use the factors that affect post-transplant survival rates in our computational experiments. The factors that have an effect on post-transplant survival are listed in Table 4.5.

4.4 COMPUTATION OF THE ORGAN REFUSAL RATES AND ARRIVAL PROBABILITIES

In this section, we calculate the overall organ refusal rates using real data and show that the decision problem described in this dissertation is faced by many liver patients. We use the UNOS3 data set to estimate the organ refusal/declination/rejection rates. Because the UNOS3 includes only the cadaveric organs that have been transplanted, the refusal rate can be calculated by simply dividing the total number of accepted offers with the total number of liver transplanted. However, as Table 1.1 shows, a significant portion of all organs are wasted. Therefore, we report the refusal rates relative to the total number of offered organs as well as the total number of cadaveric transplants. For this purpose, we estimate the number of livers that are donated between February 27, 2002 and April 30, 2003 using the UNOS web site [149]. When an organ is declined by different patients multiple times, each

Table 4.5: Factors that Affect Post-transplant Survival

Donor Characteristics	Recipient Characteristics
Gender[1]	Gender[1]
Age	Age
CMVGR[2]	CMVGR[2]
Blood type[3]	Blood type[3]
Race[4]	Mechanical ventilator[5]
	Prior transplant[6]
	Encephalopathy[7]
	Disease group[8]
	Bilirubin level
	Albumin level
	Creatinine level

[1] Binary variable that indicates whether the donor gender and the recipient gender are the same or not
[2] Binary variable that indicates whether the person has CMVGR (Cytomegalovirus) or not.
[3] Binary variable that indicates whether the donor blood type and the recipient blood type match or not
[4] Binary variable that indicates whether the donor is white or not
[5] Binary variable that indicates whether the patient is on the mechanical ventilator or not
[6] Binary variable that indicates whether the patient had another organ transplant before or not
[7] Binary variable that indicates whether the patient has encephalopathy, a situation that occurs when toxic substances accumulate in the blood or not
[8] Integer variable that shows the disease group of the patient

Table 4.6: Liver Refusal Rates by Recipient Status*

Patient Status	Refusal Rate (in %)
MELD < 11	85
MELD 11 to 18	66
MELD 19 to 24	62
MELD > 25	66

*This table reports the refusal rates with respect to the total number of cadaveric liver transplants.

refusal is counted as a separate offer and refusal. Our calculations show that 60% of all liver offers are refused by the patients/transplant surgeons. This figure rises to 68% when we consider only organs that are used for transplantation.

Table 4.6 reports the organ refusal rates by recipient MELD score with respect to the number of cadaveric liver transplantations. As can be seen from the table, patients with the lowest MELD scores have the highest refusal rates. This result can be explained intuitively as follows: the current liver allocation system implies that the sicker the patient is, the more likely it is that she will be near the "top" of the list. As a result, patients with low MELD scores typically receive lower quality organs than patients with high MELD scores. Note also that as the MELD score increases, the risk of death increases. Therefore, a patient with a low MELD score is more selective than a patient with a high MELD score. Intuitively, the refusal rate for the patients with a MELD score greater than 25 may be expected to be the lowest. However, this is not the case in Table 4.6, which may have occurred due to the sparsity of the data.

We use the UNOS2 data set to estimate $\mathcal{L}$. The estimation of $\mathcal{L}$ includes the discretization of the liver quality, the classification of the liver types and the computation of the organ arrival rates by recipient MELD score. The list of the donor characteristics that affect the post-transplant life expectancy is given in Section 4.3, which includes five donor characteristics. However, because the UNOS2 data set includes information about only the age, the race and the gender of the donor, we use these three factors to determine the quality of

the liver. We define seven categories for the donor age groups: age < 20, $20 \leq$ age < 30, $30 \leq$ age < 40, $40 \leq$ age < 50, $50 \leq$ age <60, $60 \leq$ age < 70 and age ≥ 70. The donor race consists of two categories: white or not white. The donor gender either matches the recipient's gender or does not. As a result, there are a total of 28 categories for liver types. We use the coefficients of the Cox proportional hazards model described in Section 4.3 to order the liver types. Note that the ordering of the liver types depends on the gender of the recipient. We consider only female patients throughout this study, therefore if the donor is a female then there exists a gender match, otherwise there exists a gender mismatch. Because the data are sparse, we group the liver types in two and use 14 categories for livers.

In order to find the organ arrival rates, we first compute the total number of days that each patient waits at each MELD score. Then, we calculate the total number of liver offers by recipient MELD score. Let $W_i(h)$ and $O_i(h,\ell)$ be the total number of days that patient i waits in MELD score h and the total number of type ℓ liver offers that patient i receives when her MELD score is h, respectively. Then $\mathcal{L}$ is obtained using the following formulas:

$$\mathcal{L}(L+1|h) = 1 - Min\left\{1, \frac{\sum_i \sum_\ell O_i(h,\ell)}{\sum_i W_i(h)}\right\}, h \in S_H$$

and

$$\mathcal{L}(\ell|h) = [1 - \mathcal{L}(L+1|h)]\frac{\sum_i O_i(h,\ell)}{\sum_i \sum_\ell O_i(h,\ell)}, h \in S_H \text{ and } \ell \in S_L.$$

Note that because $\mathcal{L}$ depends highly on geographical factors, we estimate a national $\mathcal{L}$ as well as 11 regional $\mathcal{L}$'s. Note also that $\frac{\sum_i \sum_\ell O_i(h,\ell)}{\sum_i W_i(h)} < 1$ holds for all $\mathcal{L}$ matrices.

5.0 LIVING-DONOR-ONLY MODEL

Much of this chapter is based on Alagoz et al. [4, 5]. The purpose of this chapter is to determine the optimal timing of living-donor liver transplantation. That is, we seek a policy describing the health states in which the living-donor liver transplantation should occur and those where waiting is the optimal action.

We assume that the donor is indifferent to the timing of transplantation and that the quality of the donated organ is fixed. We assume that there are a finite number of health states and that a complete ordering over the health states exists. We also assume that the patient is either ineligible or has decided not to receive cadaveric organ offers. Although we assume that the decision maker is both indifferent to the timing of the resolution of uncertainty and risk-neutral, we recognize that these assumptions do not necessarily hold [32]. Considering these and other types of patient preference issues is left for future research. For ease of presentation, we also assume that the disutility associated with using the living-donor liver, $\rho(h)$, is 0. Because we do not consider any cadaveric liver offers in this chapter, considering the case in which $\rho(h) > 0$ changes only the $r(h, \ell_{LD}, T_{LD})$ function and therefore does not affect the structural results in Section 5.2 but may change the control limits.

The optimal solution to this problem may appear to be straightforward, i.e. the patient should have the transplant immediately once the living donor comes forward. Consider, however, that there are two components of a patient's total reward; pre-transplant reward and post-transplant reward, and recall that the overall objective is to maximize the total discounted reward rather than maximizing one of these components. When transplantation occurs, the patient's pre-transplant life ends and the post-transplant life starts. Therefore, if transplantation occurs at an early stage of the disease, the patient may maximize her post-transplant reward but not her total reward. Figure 5.1 shows such a case. The first

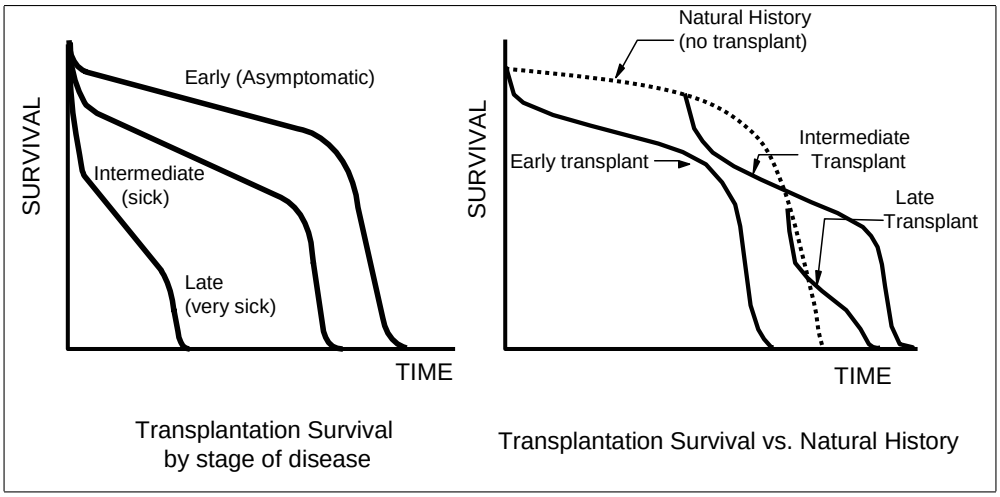

Figure 5.1: An Example Showing Why Transplanting as Early as Possible is Suboptimal

graph shows the post-transplant survival rate for the transplant recipients when they receive transplants at different stages of the liver disease. The second graph combines the post-transplant survival rate with the pre-transplant survival rate. As can be seen from the figure, although an early transplant results in the longest post-transplant survival rate, an intermediate transplant may produce a longer survival rate.

To the best of our knowledge, there are no other studies that consider the optimal timing of living-donor transplantation for any organs. The remainder of this chapter is organized as follows. In Section 5.1, we present an MDP model of the problem. We derive several structural properties of this MDP model and its optimal policy in Section 5.2. In Section 5.3, we present and discuss computational results. We draw some conclusions and discuss future research directions in Section 5.4.

5.1 MODEL FORMULATION

This problem is a special case of the model presented in Chapter 3, i.e. $\mathcal{L}(L+1|h) = 1, h \in S_H$. This is equivalent to assuming that the probability that the patient will receive a cadaveric liver offer is 0. We then can drop the liver state from the optimality equation. Note

that we will use $r(h, \ell_{LD}, T_{LD})$ throughout this chapter, because $\rho(h)$ is assumed to be 0 for all $h \in S_H$.

Figure 5.2 shows the state-transition diagram of the MDP. The decision maker can take one of two actions at state h, namely, "Transplant" or "Wait for one more decision epoch". If the patient chooses "Transplant" in health state h, she receives a reward of $r(h, \ell_{LD}, T_{LD})$, quits the process and moves to absorbing state "Transplant" with probability 1. If the patient chooses to "Wait" in health state h, then she receives an intermediate reward of $r(h, W)$ and moves to health state $h' \in S_H$ with probability $\mathcal{H}(h'|h)$. Let $V(h)$ be the maximum total expected discounted reward that the patient can attain when her current health is h. The optimal solution to this problem can be obtained by solving the following set of recursive equations [125]:

$$V(h) = \max \left\{ r(h, \ell_{LD}, T_{LD}), r(h, W) + \lambda \sum_{h' \in S_H} \mathcal{H}(h'|h) V(h') \right\}, h = 1, \ldots, H. \qquad (5.1)$$

5.2 STRUCTURAL PROPERTIES

In this section, we derive some structural properties of the living-donor-only model (LDM) given by (5.1). Some of these properties provide closed-form solutions to the LDM under certain assumptions on the reward function and the transition probability matrix. Our main result establishes sufficient conditions that ensure the existence of an optimal control-limit policy. A control-limit policy in this case is of the simple form: choose the "Transplant" action and "Accept" the organ if and only if the observed health state is one of the states $j, j+1, \ldots, H$, for the control limit j.

The following additional assumptions are made throughout:

(*As1*) The function $r(h, \ell_{LD}, T_{LD})$ is nonnegative and is nonincreasing in h. That is, as the patient gets sicker her post-transplant discounted quality-adjusted expected life days do not increase.

(*As2*) The function $r(h, W)$ is nonincreasing. That is, the expected intermediate reward that the patient accrues is nonincreasing in h.

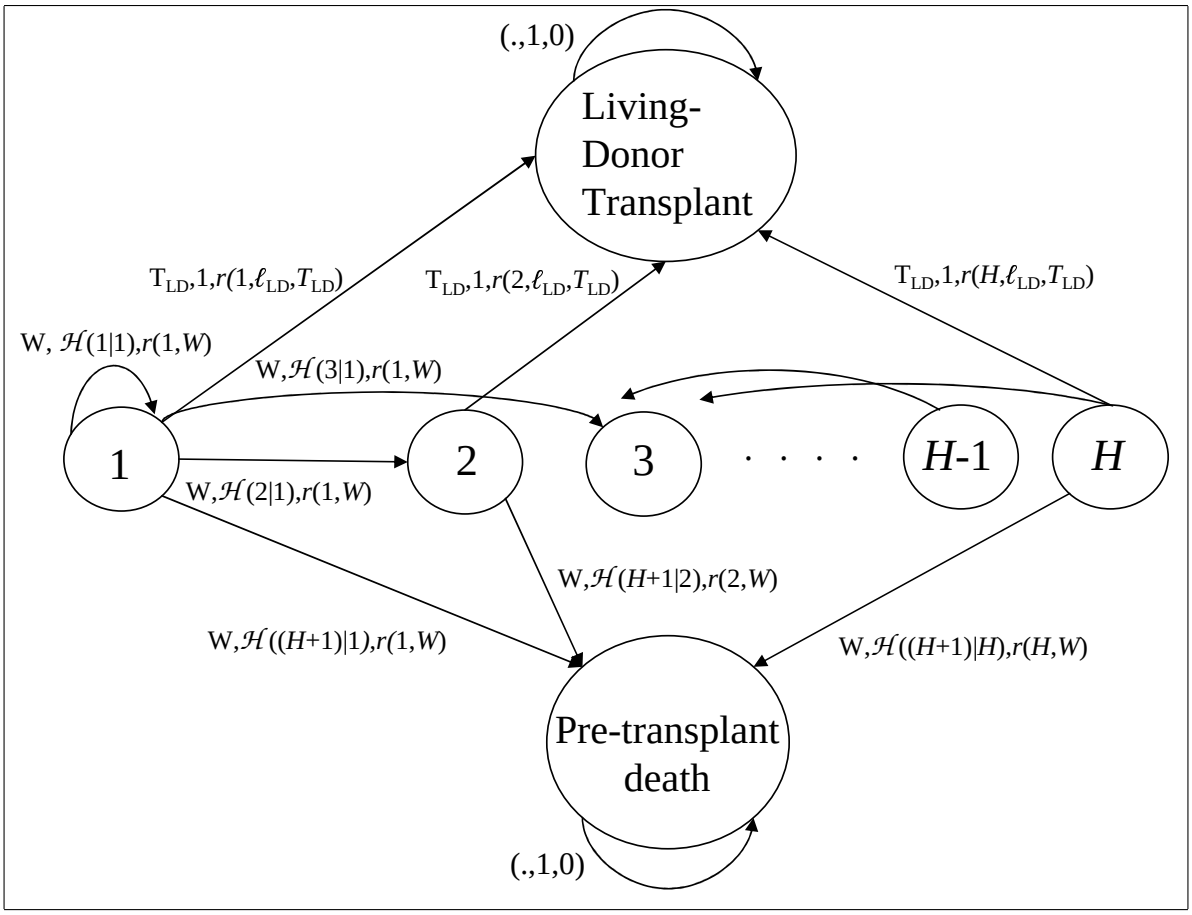

Figure 5.2: State-Transition Diagram of the Living-Donor-Only Model*

*The labels $a(h), \mathcal{H}(j|h), r(h,a)$ on each arc represent the action taken at state h, the probability that the patient will move to state j when her current state is h, and the reward obtained by taking action a in health state h, respectively. Note that this figure does not show all possible transitions.

Theorem 5.1 presents sufficient conditions under which it is optimal to choose "Transplant" for all health states.

Theorem 5.1. *If $\mathcal{H}$ is upper triangular, i.e. the patient health never improves, and*

$$[1-\lambda\mathcal{H}(h|h)]r(h,\ell_{LD},T_{LD}) \geq r(h,W)+\lambda r(h+1,\ell_{LD},T_{LD})[1-\mathcal{H}(h|h)] \text{ for } h=1,...,H, \tag{5.2}$$

then $a^*(s)$='T_{LD}' for all $s \in S_H$.

Proof. The theorem holds for $s=H$ because

$$\begin{aligned} V(H) &= \max\{r(H,\ell_{LD},T_{LD}), r(H,W)+ \\ &\quad \lambda(1-\mathcal{H}(H|H))r(H+1,\ell_{LD},T_{LD})+\lambda\mathcal{H}(H|H)V(H)\}. \end{aligned}$$

Since $r(H+1,\ell_{LD},T_{LD})=0$ and $r(H,\ell_{LD},T_{LD}) \geq \frac{r(H,W)}{1-\lambda\mathcal{H}(H|H)}$, $V(H)=r(H,\ell_{LD},T_{LD})$ and $a^*(H)$='$\mathrm{T_{LD}}$'. Assume that $a^*(s)$='$\mathrm{T_{LD}}$' for $s=h$,...,H-1. Then $V(s)=r(s,\ell_{LD},T_{LD})$ for $s=h$,...,H. Now
$V(h-1)=\max\{r(h-1,\ell_{LD},T_{LD}), r(h-1,W)+\lambda\sum_{h'}\mathcal{H}(h'|h-1)V(h')\}$, and

$$\begin{aligned} r(h-1,W)+\lambda\sum_{h'}\mathcal{H}(h'|h-1)V(h') &\leq r(h-1,W)+\lambda\sum_{h'>h-1}\mathcal{H}(h'|h-1)r(h,\ell_{LD},T_{LD})+ \\ &\quad \lambda\mathcal{H}(h-1|h-1)V(h-1) \qquad (5.3) \\ &= r(h-1,W)+\lambda r(h,\ell_{LD},T_{LD})[1-\mathcal{H}(h-1|h-1)] \\ &\quad +\lambda\mathcal{H}(h-1|h-1)V(h-1) \qquad (5.4) \end{aligned}$$

where (5.3) follows from the induction hypothesis, *(As1)* and because $\mathcal{H}$ is upper triangular. So if $a^*(h-1)$ is uniquely 'W', then (5.4) is equivalent to

$$\begin{aligned} V(h-1)[1-\lambda\mathcal{H}(h-1|h-1)] &\leq r(h-1,W)+\lambda r(h,\ell_{LD},T_{LD})[1-\mathcal{H}(h-1|h-1)] \\ &\leq r(h-1,\ell_{LD},T_{LD})[1-\lambda\mathcal{H}(h-1|h-1)], \end{aligned}$$

where the last inequality follows from (5.2). Therefore, $V(h-1) \leq r(h-1,\ell_{LD},T_{LD})$ or $a^*(h-1)$='$\mathrm{T_{LD}}$' which is a contradiction from which the result follows. □

Both the upper triangularity assumption on $\mathcal{H}$ and Condition (5.2) in Theorem 5.1 are very restrictive and unrealistic. In fact, our computational tests in Section 5.3 indicate that transplanting right away is typically suboptimal for the LDM. For this reason, we next explore conditions under which the optimal policy may not be "Transplant" for all states, but still has an appealing structure.

As discussed in Section 2.1, in many areas of application, such as maintenance optimization [29, 120, 150], inventory theory [71] and queueing [157], authors derive sufficient conditions to ensure the existence of an optimal control-limit policy. Most assume special structure on the transition probability matrix and/or the reward function. Below we define some concepts that are used to specify these special structures. Interested readers should refer to [12, 44, 45, 46] and [120] for more details.

Definition 5.1. *[12] (a) A discrete distribution* $\{p_k\}_{k=0}^{\infty}$ *is* IFR *(increasing failure rate) if* $\frac{p_k}{\sum_{i=k}^{\infty} p_i}$ *is nondecreasing in* $k = 0, 1, 2, \ldots$

(b) A Markov chain is said to be IFR if its rows are in increasing stochastic order, that is,

$$z(i) = \sum_{j=h}^{H+1} \mathcal{H}(j|i) \tag{5.5}$$

is nondecreasing in i *for all* $h = 1, \ldots, H+1$.

This definition is equivalent to the well known notion of stochastic dominance [143] and may be viewed intuitively as follows: the sicker the patient, the more probable that the patient will become even sicker.

It is interesting to note that several authors consider machine replacement problems that have similar, but not identical, structures to the LDM [12, 44, 45, 46, 130]. The main difference is that there are two absorbing states in the LDM, whereas most machine replacement problems have at most one absorbing state. As a result, standard assumptions such as IFR that guarantee the existence of a control-limit optimal policy in machine replacement problems may not be sufficient to ensure the existence of a control-limit optimal policy in our problem.

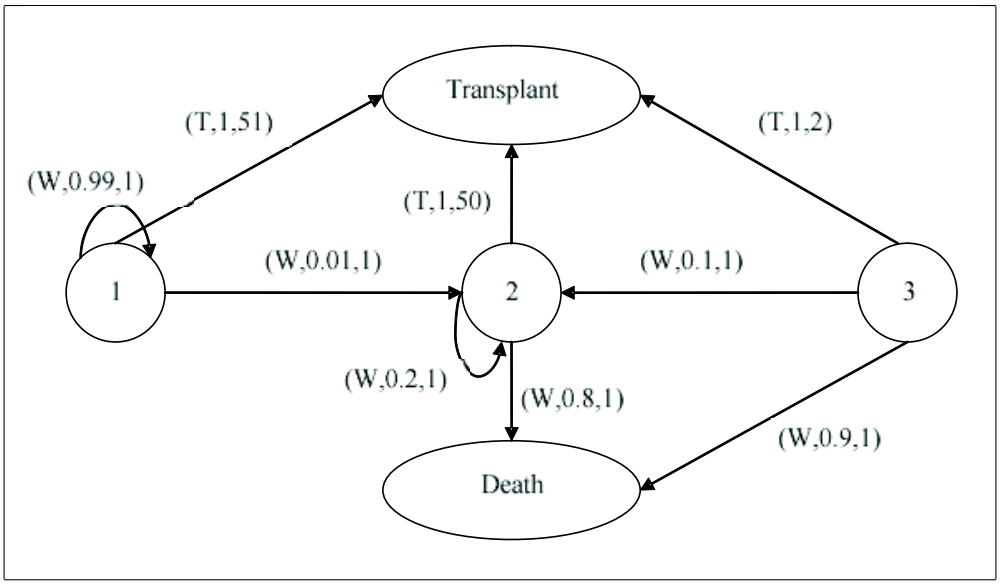

Figure 5.3: State-transition Diagram for the Counterexample

Remark 5.1. *If the transition probability matrix is IFR, an optimal control-limit policy may not exist, as demonstrated by the following counterexample.*

Consider the example in Figure 5.3. Assume $r(h, W) = 1, \forall h \in S_H$ and $\lambda = 0.99$. It is clear that $r(h, \ell_{LD}, T_{LD})$ is nonincreasing. Furthermore, the transition probability matrix is IFR based on Definition 5.1, and therefore satisfies the assumptions of the remark. However,

$V(1) = (1 + (0.01)(0.99)50)/(1 - 0.99^2) = 75.13$ and $a^*(1)$='W',

$V(2) = r(2, \ell_{LD}, T_{LD}) = 50$ and $a^*(2)$='$\mathrm{T_{LD}}$',

$V(3) = 1 + (0.99)(0.1)(50) = 5.95$ and $a^*(3)$='W'.

Although the assumptions of Remark 1 are satisfied, the optimal policy is not a control-limit policy. The result in Remark 5.1 can be explained intuitively as follows: If $r(j, \ell_{LD}, T_{LD}) - r(j+1, \ell_{LD}, T_{LD})$ is very large, then it is possible that, starting from state $j+1$, the patient's expected cumulative intermediate reward before reaching "Death" exceeds $r(j+1, \ell_{LD}, T_{LD})$, so that the "Transplant" option is not the optimal action in state $j + 1$.

Some researchers impose the assumption of *totally positive of order 2 (TP_2)* on the transition probability matrix [130]. A matrix is TP_2 if the determinants of all of its 2×2 submatrices are nonnegative. It has been shown that a TP_2 matrix is also IFR [12]. In the

above counterexample, the transition probability matrix is also TP_2, therefore even if we replace the IFR assumption with the TP_2 assumption, the remark still holds.

Before proving our main results, we give two inequalities that hold for IFR matrices and nonincreasing functions.

Lemma 5.1. *Let $\mathcal{H}$ be an IFR transition probability matrix and $V(h)$ be a nonincreasing function. Then the following hold:*

$$(a) \quad \sum_{h' \le h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)]V(h') \ge \sum_{h' \le h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)]V(h).$$

$$(b) \quad \sum_{h'' > h} [\mathcal{H}(h''|h) - \mathcal{H}(h''|h+1)]V(h'') \ge \sum_{h'' > h} [\mathcal{H}(h''|h) - \mathcal{H}(h''|h+1)]V(h+1).$$

Proof. (a) First note that the IFR assumption requires that $\sum_{i=1}^{h'} \mathcal{H}(i|h) \ge \sum_{i=1}^{h'} \mathcal{H}(i|h+1)$ for any $h' \in S_H$. Now,

$$\begin{aligned}
&\textstyle\sum_{h' \le h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)]V(h') \\
&= [\mathcal{H}(1|h) - \mathcal{H}(1|h+1)]V(1) + \sum_{h'=2}^{h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)]V(h') \\
&\ge [\mathcal{H}(1|h) - \mathcal{H}(1|h+1)]V(2) + \sum_{h'=2}^{h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)]V(h') \qquad (5.6) \\
&= [\mathcal{H}(1|h) + \mathcal{H}(2|h) - \mathcal{H}(1|h+1) - \mathcal{H}(2|h+1)]V(2) + \sum_{h'=3}^{h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)]V(h') \\
&\ge [\mathcal{H}(1|h) + \mathcal{H}(2|h) - \mathcal{H}(1|h+1) - \mathcal{H}(2|h+1)]V(3) + \sum_{h'=3}^{h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)]V(h')
\end{aligned}$$

where (5.6) follows because $\mathcal{H}(1|h) \ge \mathcal{H}(1|h+1)$ and $V(1) \ge V(2)$. We obtain the second inequality by recasting (5.6), and the last inequality holds because $\mathcal{H}(1|h) + \mathcal{H}(2|h) \ge \mathcal{H}(1|h+1) + \mathcal{H}(2|h+1)$ and $V(2) \ge V(3)$. $\mathcal{H}(1|h+1) + \mathcal{H}(2|h+1)$ and $V(2) \ge V(3)$. The result follows if we apply the same procedure to states 3 through h.

(b) The proof is similar to the proof of Part (a) and is omitted. □

In the following theorem, we consider the case in which the transition probability matrix is IFR and show that the patient's total discounted expected life days is nonincreasing in h.

Theorem 5.2. *If $\mathcal{H}$ is IFR, then for $s = 1, ..., H$, $V(s) \ge V(s+1)$.*

Proof. This theorem is a direct result of the infinite-horizon version of Lemma 3.9.4 in Topkis [143] or, equivalently Theorem 4.7.3 in Puterman [125]. We also provide an alternative proof below.

We show that if we apply the value iteration algorithm to solve this problem, at any step of the algorithm the monotonicity of the value function is preserved. Since it is well known that the value iteration algorithm converges [125], the result will then follow.

Let $V^i(h)$ be the value associated with state h at the ith iteration of the value iteration algorithm. Assume that the value iteration algorithm starts with a value of 0 for each health state, i.e. $V^0(h) = 0$ for all $h \in S_H$. For the base case, we need to prove that $V^1(h) \geq V^1(h+1)$, for $h = 1, \ldots, H$, where $V^1(h) = \max\left\{r(h, \ell_{LD}, T_{LD}), r(h, W) + \lambda \sum_{h'} \mathcal{H}(h'|h) V^0(h')\right\}$. Because $V^0(h') = 0$ for $h' \in S_H$, $V^1(h) = \max\{r(h, \ell_{LD}, T_{LD}), r(h, W)\}$. Because both $r(h, \ell_{LD}, T_{LD})$ and $r(h, W)$ is monotonically nonincreasing, it is obvious that the value function is also monotonically nonincreasing.

Now, suppose that the property holds for iterations $i = 2, \ldots, n$, that is, $V^i(h) \geq V^i(h+1)$ for $h = 1, \ldots, H$ and $i = 2, \ldots, n$. Note that

$$V^{n+1}(h) = \max\left\{r(h, \ell_{LD}, T_{LD}), r(h, W) + \lambda \sum_{h'} \mathcal{H}(h'|h) V^n(h')\right\} \quad \text{and}$$

$$V^{n+1}(h+1) = \max\left\{r(h+1, \ell_{LD}, T_{LD}), r(h+1, W) + \lambda \sum_{h'} \mathcal{H}(h'|h+1) V^n(h')\right\}.$$

If $V^{n+1}(h+1) = r(h+1, \ell_{LD}, T_{LD})$, $V^{n+1}(h) \geq r(h, \ell_{LD}, T_{LD}) \geq r(h+1, \ell_{LD}, T_{LD}) = V^{n+1}(h+1)$ and the result immediately follows. Otherwise, we have the following:

$$\begin{aligned}
& V^{n+1}(h) - V^{n+1}(h+1) \\
\geq\ & r(h, W) - r(h+1, W) + \lambda \left[\sum_{h' \leq h} \mathcal{H}(h'|h) V^n(h') + \sum_{h'' > h} \mathcal{H}(h''|h) V^n(h'')\right] - \\
& \lambda \left[\sum_{h' \leq h} \mathcal{H}(h'|h+1) V^n(h') + \sum_{h'' > h} \mathcal{H}(h''|h+1) V^n(h'')\right] \\
\geq\ & \lambda \left(\sum_{h' \leq h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)] V^n(h') + \right. \\
& \left. \sum_{h'' > h} [\mathcal{H}(h''|h) - \mathcal{H}(h''|h+1)][\mathcal{H}(h''|h) - \mathcal{H}(h''|h+1)] V^n(h'')\right), \qquad (5.7)
\end{aligned}$$

where (5.7) follows because of (As2). From the induction assumption, $V^n(h)$ is monotonic. Therefore, as a result of Lemma 5.1, the inequality is not violated when we replace $V^n(h')$ with $V^n(h)$ for $h' \leq h$ and $V^n(h'')$ with $V^n(h+1)$ for $h'' > h$. Hence, the above inequality can be rewritten as follows:

$$
\begin{aligned}
& V^{n+1}(h) - V^{n+1}(h+1) \\
& \geq \lambda \left(\sum_{h' \leq h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)] V^n(h) + \sum_{h'' > h} [\mathcal{H}(h''|h) - \mathcal{H}(h''|h+1)] V^n(h+1) \right) \\
& = \lambda \left(\sum_{h' \leq h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)] [V^n(h) - V^n(h+1)] \right).
\end{aligned}
$$

By the IFR property and the induction assumption, both $V^n(h) - V^n(h+1)$ and $\sum_{h' \leq h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)]$ are nonnegative. Therefore, $V^{n+1}(h) - V^{n+1}(h+1)$ is also nonnegative, from which the result follows. □

The following theorem, the main result of this section, gives a set of sufficient conditions to guarantee the existence of a control-limit optimal policy.

Theorem 5.3. *Let $\mathcal{H}$ be an IFR matrix and suppose $\mathcal{H}$ and $r(h, \ell_{LD}, T_{LD})$ satisfy the following conditions:*

$$
\sum_{k=j}^{H} \mathcal{H}(k|h) \leq \sum_{k=j}^{H} \mathcal{H}(k|h+1), \text{ for } j = h+1, \ldots, H \text{ and } h = 1, \ldots, H, \quad \text{and} \tag{5.8}
$$

$$
\frac{r(h, \ell_{LD}, T_{LD}) - r(h+1, \ell_{LD}, T_{LD})}{r(h, \ell_{LD}, T_{LD})} \leq \lambda [\mathcal{H}(H+1|h+1) - \mathcal{H}(H+1|h)], \quad \text{for } h = 1, \ldots, H-1. \tag{5.9}
$$

Then there exists an optimal control-limit policy. In other words there exists a state j such that $a^*(1) = a^*(2) = \ldots = a^*(j-1)$='$W$' and $a^*(j) = a^*(j+1) = \ldots = a^*(H)$='$T_{LD}$'.

Proof. First note that the monotonicity result of Theorem 5.2 holds. If we show that for some h, $a^*(h)$='$\mathrm{T_{LD}}$' implies $a^*(h+1)$='$\mathrm{T_{LD}}$' then the result follows. Assume that the converse is true. In other words, for some h, $a^*(h)$='$\mathrm{T_{LD}}$' but $a^*(h+1)$ is uniquely 'W'. In this case, $r(h,\ell_{LD},T_{LD}) \geq r(h,W) + \lambda \sum_{h' \in S_H} \mathcal{H}(h'|h)V(h')$ and $r(h+1,\ell_{LD},T_{LD}) < r(h+1,W) + \lambda \sum_{h' \in S_H} \mathcal{H}(h'|h+1)V(h')$. It is obvious that,

$$\begin{aligned} r(h,\ell_{LD},T_{LD}) - r(h+1,\ell_{LD},T_{LD}) \;>\; & r(h,W) - r(h+1,W) + \\ & \lambda \left[\sum_{h' \leq h} \mathcal{H}(h'|h)V(h') + \sum_{h''>h} \mathcal{H}(h''|h)V(h'') \right] - \\ & \lambda \left[\sum_{h' \leq h} \mathcal{H}(h'|h+1)V(h') + \sum_{h''>h} \mathcal{H}(h''|h+1)V(h'') \right], \\ \geq \; & \lambda \left[\sum_{h' \leq h} \mathcal{H}(h'|h)V(h') + \sum_{h''=h+1}^{H} \mathcal{H}(h''|h)V(h'') \right] - \\ & \lambda \left[\sum_{h' \leq h} \mathcal{H}(h'|h+1)V(h') + \sum_{h''=h+1}^{H} \mathcal{H}(h''|h+1)V(h'') \right], \end{aligned}$$

because of *(As2)* and $V(H+1) = 0$. The last inequality can also be rewritten as

$$\begin{aligned} r(h,\ell_{LD},T_{LD}) - r(h+1,\ell_{LD},T_{LD}) > \lambda \Bigg(& \sum_{h' \leq h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)]V(h') + \\ & \sum_{h''=h+1}^{H} [\mathcal{H}(h''|h) - \mathcal{H}(h''|h+1)]V(h'') \Bigg). \end{aligned} \tag{5.10}$$

From Theorem 5.2 we know that $V(h)$ is nonincreasing in h. Therefore, the result of Lemma 5.1 applies to this problem and we can replace each $V(h')$ with $V(h)$ without violating (5.10). Similarly, as a result of Lemma 5.1 and (5.8), we can replace each $V(h'')$ with $V(h+1)$. We then obtain the following:

$$
\begin{aligned}
& r(h,\ell_{LD},T_{LD}) - r(h+1,\ell_{LD},T_{LD}) \\
& \quad > \lambda\Bigg(\Bigg[\Bigg(1-\sum_{h''=h+1}^{H}\mathcal{H}(h''|h)-\mathcal{H}(H+1|h)\Bigg)- \\
& \qquad \Bigg(1-\sum_{h''=h+1}^{H}\mathcal{H}(h''|h+1)-\mathcal{H}(H+1|h+1)\Bigg)\Bigg]V(h) \\
& \qquad +\sum_{h''=h+1}^{H}[\mathcal{H}(h''|h)-\mathcal{H}(h''|h+1)]V(h+1)\Bigg), \\
& \quad = \lambda\Bigg(\sum_{h''=h+1}^{H}[\mathcal{H}(h''|h+1)-\mathcal{H}(h''|h)][V(h)-V(h+1)]\Bigg)+ \\
& \qquad \lambda\left([\mathcal{H}(H+1|h+1)-\mathcal{H}(H+1|h)]V(h)\right).
\end{aligned}
$$

From (5.8), $\sum_{h''=h+1}^{H}[\mathcal{H}(h''|h+1)-\mathcal{H}(h''|h)]$ is nonnegative, so we can drop the first term and rewrite the last inequality as follows:

$$
r(h,\ell_{LD},T_{LD}) - r(h+1,\ell_{LD},T_{LD}) > \lambda[\mathcal{H}(H+1|h+1)-\mathcal{H}(H+1|h)]V(h). \tag{5.11}
$$

Using (5.11) and (5.9) we obtain $\lambda[\mathcal{H}(H+1|h+1)-\mathcal{H}(H+1|h)]V(h) < \lambda[\mathcal{H}(H+1|h+1)-\mathcal{H}(H+1|h)]r(h,\ell_{LD},T_{LD})$, which is equivalent to assuming that $V(h) < r(h,\ell_{LD},T_{LD})$ which means there exists a contradiction. Therefore, $a^*(h+1)$='$\mathrm{T_{LD}}$' also holds, from which the result follows. □

In Theorem 5.3, (5.8) implies that the sicker the patient is, the more likely it is that she will move to sicker health states. Note that (5.8) and Definition 5.1 have similar interpretations, but (5.8) is neither a consequence of, nor sufficient to establish Definition 5.1. Condition (5.9) on the reward function has an intuitive explanation. Namely, as the patient gets sicker, the reduction in the benefit of waiting is greater than the reduction in the benefit of performing the transplant.

In the following set of theorems, we perform sensitivity analysis on the LDM, i.e., we derive sufficient conditions for situations in which the optimal policy is the same. Theorem 5.4 considers two problems which have equivalent transition probabilities and discount factors, but different reward functions. Given a set of states for which the optimal decision is to "Transplant" in the first problem, we present conditions under which "Transplant" remains

optimal for the same set of states in the second problem. This theorem reflects the real-life situation in which a patient has several potential living donors with different organ qualities. The number of health states for which it is optimal to "Transplant", the *"Transplant" region*, decreases as the organ quality decreases. The computational results in Section 5.3 confirm this theoretical result.

Theorem 5.4. *Let* $r_2(h, \ell_{LD}, T_{LD}) = (1+\alpha) r_1(h, \ell_{LD}, T_{LD}), \forall h \in S_H$ *where* $r_1(h, \ell_{LD}, T_{LD})$ *and* $r_2(h, \ell_{LD}, T_{LD})$ *are two different reward functions. Let* $a_1^*(h)$ *and* $a_2^*(h)$ *be the optimal policies when the reward functions are* $r_1(h, \ell_{LD}, T_{LD})$ *and* $r_2(h, \ell_{LD}, T_{LD})$*, respectively. Then the following are true:*

(a) If $\alpha > 0$, $a_1^*(h)$='T_{LD}' *implies that* $a_2^*(h)$='T_{LD}'.

(b) If $\alpha < 0$, $a_1^*(h)$='W' *implies that* $a_2^*(h)$='W'.

Proof. (a) Let $V_i(h)$ be the optimal value for state h for problem i, i.e., the problem having the reward function $r_i(h, \ell_{LD}, T_{LD})$,$i = 1, 2$.

Under the above definitions, we first show that at any step n of the value iteration, if $V_2^n(h') \leq V_1(h')(1+\alpha)$ for all $h' \in S_H$, then $V_2^{n+1}(h) \leq V_1(h)(1+\alpha)$ holds for any state h.

If $r_2(h, \ell_{LD}, T_{LD}) \geq r(h, W) + \sum_{h' \in S_H} \lambda \mathcal{H}(h'|h) V_2^n(h')$ then $V_2^{n+1}(h) = r_2(h, \ell_{LD}, T_{LD}) = r_1(h, \ell_{LD}, T_{LD})(1+\alpha) \leq V_1(h)(1+\alpha)$ follows immediately.

Otherwise, the proof proceeds as follows:

$$
\begin{aligned}
r(h, W) + \lambda \sum_{h' \in S_H} \mathcal{H}(h'|h) V_2^n(h') &\leq r(h, W) + \lambda \sum_{h' \in S_H} \mathcal{H}(h'|h) V_1(h')(1+\alpha) \quad (5.12) \\
&= \left(1 + \lambda \sum_{h' \in S_H} \mathcal{H}(h'|h) V_1(h')\right) + \\
&\quad \alpha \left(\lambda \sum_{h' \in S_H} \mathcal{H}(h'|h) V_1(h')\right) \\
&\leq V_1(h) + \alpha V_1(h) = V_1(h)(1+\alpha), \quad (5.13)
\end{aligned}
$$

where both (5.12) and (5.13) follow from the definition of the optimal value function.

Now suppose that we use the value iteration algorithm for solving the second problem with the following initial values

$$V_2^0(h) = r_2(h, \ell_{LD}, T_{LD}).$$

For any $h \in S_H$, at each iteration i of the value iteration algorithm, $V_2^i(h) \leq V_1(h)(1+\alpha)$ holds as shown in 5.13. Therefore, if $a_1^*(h)$='$\mathrm{T_{LD}}$' for some h, then

$$V_2(h) \leq V_1(h)(1+\alpha) = r_1(h, \ell_{LD}, T_{LD})(1+\alpha) \leq r_2(h, \ell_{LD}, T_{LD}),$$

so $a_2^*(h)$='$\mathrm{T_{LD}}$' also holds.

(b) The proof of Part (b) is similar to the proof of Part (a) and is omitted. □

Our computational tests show that as the discount rate increases, the number of health states that it is optimal to "Wait", *the optimal "Wait" region*, also increases. An intuitive explanation for this result is given in Section 5.3. The following theorem presents a set of sufficient conditions that guarantees the enlargement in the "Wait" region as the discount rate increases. In this theorem we consider expected life days rather than quality-adjusted life days. Let $ELD(h)$ be the total expected post-transplant life days, given that the transplant is performed when the patient health is h. Note that $r(h, \ell_{LD}, T_{LD})$ is in fact a function of $ELD(h)$ and λ, namely, $r(h, \ell_{LD}, T_{LD}) = \frac{1-\lambda^{ELD(h)}}{1-\lambda}$. In our previous discussion, we suppressed $r(h, \ell_{LD}, T_{LD})$'s dependence on λ and $ELD(h)$ for ease of notation. However, in the next theorem, since the discount rate affects $r(h, \ell_{LD}, T_{LD})$, we give the sufficient conditions in terms of $ELD(h)$ and λ.

Theorem 5.5. *For any state h with $a^*(h)$='W', if $ELD(h)\lambda^{ELD(h)-1}(1-\lambda\mathcal{H}(h|h)) \geq r(h, W)$ for $0 \leq \underline{\lambda} \leq \lambda \leq \overline{\lambda} < 1$ for an arbitrary interval $[\underline{\lambda}, \overline{\lambda}]$, then as λ increases within the interval $[\underline{\lambda}, \overline{\lambda}]$ the optimal action for state h does not change.*

Proof. If we show that the difference between the value of the "Wait" action and the "Transplant" action does not decrease as λ increases then, the result follows. Therefore we need to prove that the function

$$g(\lambda) = \frac{r(h, W) + \lambda \sum_{h' \in S_H, h' \neq h} \mathcal{H}(h'|h) V(h')}{1 - \lambda \mathcal{H}(h|h)} - \frac{1-\lambda^{ELD(h)}}{1-\lambda} \qquad \text{is increasing in } \lambda.$$

It can easily be shown that as λ increases, $V(h')$ also increases for all $h' \in S_H$. Since $\lambda \sum_{h' \in S_H, h' \neq h} \mathcal{H}(h'|h)V(h')$ is nonnegative, we can drop this term and show that the remaining function increases as λ increases. If we rearrange the terms we have the following

$$g'(\lambda) = \frac{r(h,W)[1-\lambda] - (1-\lambda^{ELD(h)})(1-\lambda\mathcal{H}(h|h))}{(1-\lambda\mathcal{H}(h|h))(1-\lambda)},$$

which decreases in the denominator as λ increases. Therefore, it is sufficient to show that the value of the numerator increases as λ increases. In order to show that we take derivative with respect to λ and obtain the following:

$$\begin{aligned}
\frac{\partial g'(\lambda)}{\partial \lambda} &= \frac{\partial [r(h,W)[1-\lambda] - (1-\lambda^{ELD(h)})(1-\lambda\mathcal{H}(h|h))]}{\partial \lambda} \\
&= \frac{\partial [r(h,W)[1-\lambda] - 1 + \lambda\mathcal{H}(h|h) + \lambda^{ELD(h)} - \mathcal{H}(h|h)\lambda^{ELD(h)+1}]}{\partial \lambda} \\
&= \mathcal{H}(h|h) - r(h,W) + ELD(h)\lambda^{ELD(h)-1} - \mathcal{H}(h|h)(ELD(h)+1)\lambda^{ELD(h)} \\
&= \mathcal{H}(h|h) - \mathcal{H}(h|h)\lambda^{ELD(h)} + ELD(h)\lambda^{ELD(h)-1} - \\
&\quad \mathcal{H}(h|h)ELD(h)\lambda^{ELD(h)} - r(h,W)
\end{aligned}$$

in which $\mathcal{H}(h|h) - \mathcal{H}(h|h)\lambda^{ELD(h)} \geq 0$ and the remaining term can be rearranged as follows:

$$ELD(h)\lambda^{ELD(h)-1}(1-\lambda\mathcal{H}(h|h)) - r(h,W),$$

which is ≥ 0 when $\lambda \in [\underline{\lambda}, \overline{\lambda}]$ from the assumption. Therefore the derivative is nonnegative which shows that as λ increases in the range of $[\underline{\lambda}, \overline{\lambda}]$ the value of the function also increases from which the result follows. □

In Theorem 5.5, the condition $ELD(h)\lambda^{ELD(h)-1}(1-\lambda\mathcal{H}(h|h)) \geq r(h,W)$ appears to be restrictive. However, our computational tests show that if $r(h,W)$ equals 1, all of our clinical data with an annual discount rate greater than 0.5 satisfy the condition.

Theorem 5.6 addresses the relationship between the optimal policies for two patients with different disease progression rates but equivalent reward functions. We show that, given both patients have control-limit optimal policies, if one of the patients deteriorates faster than the other, then she has a lower optimal control limit. First, we describe a plausible relationship between the transition probability matrices for two different diseases. Namely, if one of the transition probability matrices has a faster deterioration rate than another, then we say that

transition probability matrix is dominated by the other. This condition may arise when the progression of two diseases are different.

Definition 5.2. *Let* $P = [P(j|i)], i, j = 1, \ldots, n$ *and* $Q = [Q(j|i)], i, j = 1, \ldots, n$ *be two transition probability matrices. We say* P *dominates* Q, $P \succeq Q$, *if* $\sum_{j=k}^{n} P(j|i) \leq \sum_{j=k}^{n} Q(j|i), \quad 1 \leq i, k \leq n.$

Consider the implications of this definition by letting the random variables $X(h)$ and $Y(h)$ be the time to death starting from health state h under transition probability matrices P and Q if no transplant is performed, respectively. It can easily be shown that Definition 5.2 implies ordinary stochastic dominance between $X(h)$ and $Y(h)$, i.e. $X(h) \geq_{st} Y(h)$ for $h \in S_H$.

Lemma 5.2. *Let* P *and* Q *be* $n \times n$ *transition probability matrices where* $P \succeq Q$. *Furthermore, let* $V(j)$ *be a monotonically nonincreasing function in* j. *Then for any* i, *the following are true:*

(a) $\sum_{j \leq i}[P(j|i) - Q(j|i)]V(j) \geq \sum_{j \leq i}[P(j|i) - Q(j|i)]V(i)$.

(b) $\sum_{j > i}[P(j|i) - Q(j|i)]V(j) \geq \sum_{j > i}[P(j|i) - Q(j|i)]V(i+1)$.

Proof. The proof of this lemma is very similar to that of Lemma 5.1 and is omitted. □

Theorem 5.6. *Let* Π_1 *and* Π_2 *be two instances that satisfy the conditions of Theorem 5.3 so that the optimal policies for* Π_1 *and* Π_2 *are both control-limit optimal policies with control limits* j_1 *and* j_2, *respectively. Let* $\mathcal{H}_1$ *and* $\mathcal{H}_2$ *be the transition probability matrices of* Π_1 *and* Π_2, *respectively. If* Π_1 *and* Π_2 *have the same reward functions,* $r(h, \ell_{LD}, T_{LD})$ *and* $r(h, W)$ *and* $\mathcal{H}_1 \succeq \mathcal{H}_2$, *then* $j_1 \geq j_2$.

Proof. Suppose that we solve the two problems simultaneously using the value iteration algorithm. We first show that starting with a value of 0 for all states in both problems, at the end of each iteration of the algorithm, the value function of Π_1 will be greater than or equal to the value function of Π_2 for each health state. Let $V_i^j(h)$ be the value function of the state h of problem i at the end of iteration j. We start with 0 for both problems. In this case, $V_1^1(h) = V_2^1(h) = \max\{r(h, \ell_{LD}, T_{LD}), r(h, W)\}, h = 1, \ldots, H$ as shown in Theorem 5.2. Therefore the result holds for the base case.

Now, assume that $V_1^n(h) \geq V_2^n(h),\ h = 1, \ldots, H$, holds for iterations $2, \ldots, n$. Then we want to show that $V_1^{n+1}(h) \geq V_2^{n+1}(h),\ h = 1, \ldots, H$. If for any state h, $V_2^{n+1}(h) = r(h, \ell_{LD}, T_{LD})$, then the result immediately follows since $V_1^{n+1}(h) \geq r(h, \ell_{LD}, T_{LD})$. Otherwise, the application of the value iteration algorithm results in the following:

$$V_1^{n+1}(h) \geq r(h, W) + \lambda \sum_{h' \leq h} \mathcal{H}_1(h'|h) V_1^n(h') + \lambda \sum_{h'' > h} \mathcal{H}_1(h''|h) V_1^n(h'')$$

and

$$V_2^{n+1}(h) = r(h, W) + \lambda \sum_{h' \leq h} \mathcal{H}_2(h'|h) V_2^n(h') + \lambda \sum_{h'' > h} \mathcal{H}_2(h''|h) V_2^n(h'').$$

We easily obtain the following:

$$\begin{aligned}
& V_1^{n+1}(h) - V_2^{n+1}(h) \\
\geq\ & \lambda \sum_{h' \leq h} \mathcal{H}_1(h'|h) V_1^n(h') + \lambda \sum_{h'' > h} \mathcal{H}_1(h''|h) V_1^n(h'') - \\
& \lambda \sum_{h' \leq h} \mathcal{H}_2(h'|h) V_2^n(h') - \lambda \sum_{h'' > h} \mathcal{H}_2(h''|h) V_2^n(h'') \\
\geq\ & \lambda \sum_{h' \leq h} \mathcal{H}_1(h'|h) V_2^n(h') + \lambda \sum_{h'' > h} \mathcal{H}_1(h''|h) V_2^n(h'') - \\
& \lambda \sum_{h' \leq h} \mathcal{H}_2(h'|h) V_2^n(h') - \lambda \sum_{h'' > h} \mathcal{H}_2(h''|h) V_2^n(h'') \qquad (5.14) \\
=\ & \lambda \left(\sum_{h' \leq h} [\mathcal{H}_1(h'|h) - \mathcal{H}_2(h'|h)] V_2^n(h') + \sum_{h'' > h} [\mathcal{H}_1(h''|h) - \mathcal{H}_2(h''|h)] V_2^n(h'') \right) \qquad (5.15) \\
\geq\ & \lambda \left(\sum_{h' \leq h} [\mathcal{H}_1(h'|h) - \mathcal{H}_2(h'|h)] V_2^n(h) + \sum_{h'' > h} [\mathcal{H}_1(h''|h) - \mathcal{H}_2(h''|h)] V_2^n(h+1) \right) \qquad (5.16) \\
=\ & \lambda \left(\sum_{h' \leq h} [\mathcal{H}_1(h'|h) - \mathcal{H}_2(h'|h)] [V_2^n(h) - V_2^n(h+1)] \right) \geq 0 \qquad (5.17)
\end{aligned}$$

where (5.14) follows from the induction assumption and (5.15) is obtained by simply rearranging terms. Inequality (5.16) holds because $\mathcal{H}_1 \succeq \mathcal{H}_2$ and the monotonicity of the value function imply that $V_2^n(h')$ can be replaced with $V_2^n(h)$, and $V_2^n(h'')$ can be replaced with $V_2^n(h+1)$, without violating the inequality as a result of Lemma 5.2. The first inequality in (5.17) follows from rearranging the terms in (5.16); and the second part of (5.17) follows because $\mathcal{H}_1 \succeq \mathcal{H}_2$ and the value function is monotonic.

Since the value function for Π_1 is always greater than or equal to that of Π_2 at each iteration of the value iteration algorithm, the optimal value function of Π_1 will always be greater than or equal to that of Π_2. Hence, if for state j_1, $a_1^*(j_1)$='$\mathrm{T_{LD}}$' in the first problem, since $r(j_1, \ell_{LD}, T_{LD}) \leq V_2(j_1) \leq V_1(j_1) = r(j_1, \ell_{LD}, T_{LD})$, $a_2^*(j_1)$='$\mathrm{T_{LD}}$' always holds in the second problem. As a result, since we have a control-limit optimal policy in both problems, the control limit in the first problem will always be greater than or equal to the control limit in the second problem, from which the result follows. □

Note that Theorem 5.6 holds even if the conditions of Theorem 5.3 do not hold, as long as both problems have control-limit optimal policies and nonincreasing optimal value functions and the same reward functions.

Corollary 5.1. *Let Π_1 and Π_2 be two problems having nonincreasing optimal value functions and the optimal policies for Π_1 and Π_2 are both control-limit optimal policies with control limits j_1 and j_2. Let $\mathcal{H}_1$ and $\mathcal{H}_2$ be the transition probability matrices of Π_1 and Π_2, respectively. If Π_1 and Π_2 have the same reward functions, $r(h, \ell_{LD}, T_{LD})$ and $r(h, W)$ and $\mathcal{H}_1 \succeq \mathcal{H}_2$ then $j_1 \geq j_2$.*

5.3 COMPUTATIONAL RESULTS

5.3.1 Estimating Parameters

It is well known that MDPs are subject to the so called "curse of dimensionality" [13, 33]. That is, it becomes computationally difficult to solve the problem as the problem size increases. As a result, it may be hard to solve some of the models in this dissertation using very specific health states. Furthermore, the accuracy of parameter estimations might be lost as the state space grows because some of the data are very sparse. Therefore, the states must be defined with these trade-offs between medical accuracy, data availability and computational time in mind. We use the MELD scores to represent patient health. MELD scores are restricted to integers ranging from 6 to 40, where 40 is the sickest. In this study, we use this same range to measure a patient's health status, but due to the sparsity of the

data, we aggregate the scores into groups of two or three. We do not include Status 1 as part of health state space because there are very few Status 1 patients. For instance, as of May 14, 2004, there were 15 Status 1 patients nationally whereas there were a total of 17,400 liver patients [149].

The data are classified into five disease groups based on the underlying etiology of ESLD. We estimate the transition probabilities for each disease group separately because the progression of liver disease is highly disease-dependent [47, 121]. For each group, we estimate the pre-transplant transition probabilities between health states using the NHM described in Section 4.2.

There are two types of rewards in the MDP model: pre-transplant and post-transplant. Because we are unaware of any existing data on quality-adjusted rewards for MELD scores, we use total discounted life expectancy in days rather than total discounted QALD for $r(h, W)$ and $r(h, \ell_{LD}, T_{LD})$ in our computational tests.

If the patient chooses to "Wait", the patient accrues one day as the intermediate reward, i.e. $r(h, W) = 1, \forall h \in S_H$. If the patient chooses the "Transplant" option, then she receives a post-transplant reward which is equal to the expected life days of the patient given her health status at the time of the transplant and the liver quality. To estimate the expected post-transplant life days of the patient given her MELD score at the time of transplant and liver quality, we use the Post-transplant Survival Model described in Section 4.3. The inputs of this model include the donor characteristics and the clinical characteristics of the patient at the time of the transplant.

5.3.2 Numerical Examples

First consider a 60-year-old female patient with primary biliary cirrhosis who has blood type A. Figure 5.4 depicts the optimal transplant/wait policy for this patient with ten potential donors. Table 5.2 contains the characteristics of these donors. We used these characteristics because the model of [127] identifies these factors to be most influential on the post-transplant survival rates. We apply a 0.99 annual discount rate (daily $\lambda = 0.999972$). Recall that as the MELD score increases, the patient gets sicker. In Figure 5.4, the organs are ordered from

Table 5.1: Maximum Violations of Conditions of Theorem 5.3

	Disease 1	Disease 2	Disease 3
ϵ_1	0.0234	0.0021	0.0189
ϵ_2	0.0038	0.0012	0.0189
ϵ_3	0.1052	0.1089	0.1195

lowest to highest quality, i.e. Organ 1 is the "best" and Organ 10 is the "worst" organ. The quality of a liver for a given patient is determined by the post-transplant survival model of [127]. As Figure 5.4 shows, the optimal policy varies for different liver qualities. If the only liver available to the patient is Organ 1, then the optimal policy for this particular patient is as follows: "Wait" until the MELD score rises over 14 and then accept the liver and have the transplantation.

We tested a total of 840 instances, and, while none strictly satisfied the conditions of Theorem 5.3, all of the optimal policies were of control-limit type. To quantify the magnitude of the violations of the conditions of Theorem 5.3 we define the following metrics:

$\epsilon_1 = \max_{j,h} \left\{ \max \left\{ 0, \sum_{k=j}^{H+1} [\mathcal{H}(k|h+1) - \mathcal{H}(k|h)] \right\} \right\}$ for $j = 1, \ldots, H+1$ and $h = 1, \ldots, H-1$, the maximum violation of the IFR assumption.

$\epsilon_2 = \max_{j,h} \left\{ \max \left\{ 0, \sum_{k=j}^{H} [\mathcal{H}(k|h+1) - \mathcal{H}(k|h)] \right\} \right\}$ for $j = h+1, \ldots, H$ and $h = 1, \ldots, H-1$, the maximum violation of Condition (5.8).

$\epsilon_3 = \max_h \left\{ \max \left\{ 0, \lambda \max\{0, \mathcal{H}(H+1|h+1) - \mathcal{H}(H+1|h)\} - \frac{[r(h,\ell_{LD},T_{LD}) - r(h+1,\ell_{LD},T_{LD})]}{r(h,\ell_{LD},T_{LD})} \right\} \right\}$ for $h = 1, \ldots, H-1$, the maximum violation of Condition (5.9).

Note that the values of ϵ_1 and ϵ_2 depend only on the etiology of the liver disease whereas ϵ_3 depends on the patient type, donor organ quality and the discount rate because $r(h, \ell_{LD}, T_{LD})$ is a function of these three factors. Table 5.1 reports the maximum ϵ_1, ϵ_2 and ϵ_3 values obtained with an annual discount rate of 0.99. The values for ϵ_1 and ϵ_2 are very small whereas ϵ_3 values are relatively larger.

As discussed earlier, the policies in Figure 5.4 are examples of control-limit policies. The control limit for this patient is a MELD score of 16 when the potential donors are of organ

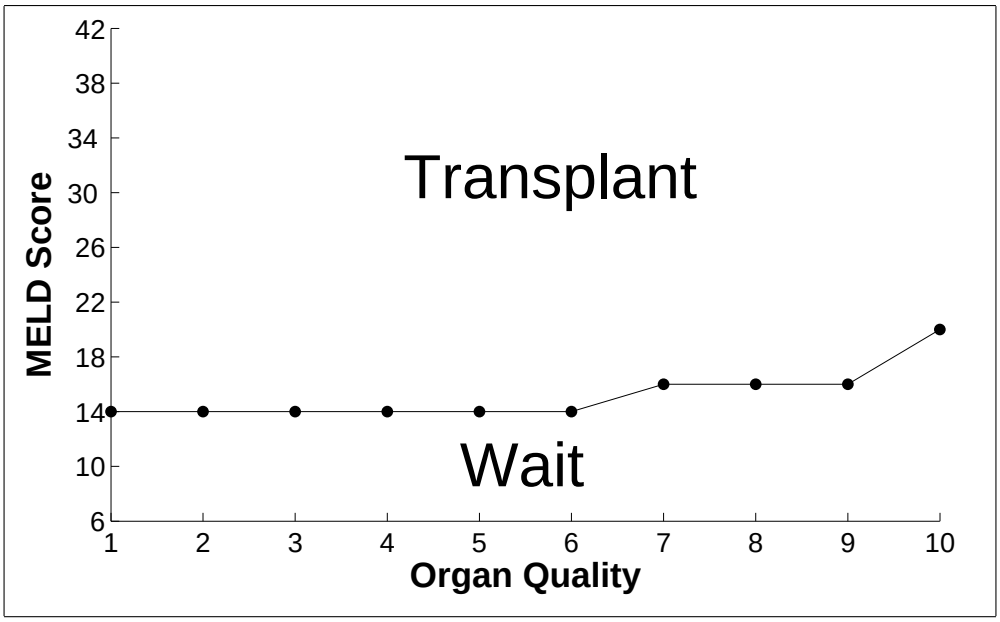

Figure 5.4: Transplant-Wait Decisions for Annual Discount Rate=0.99 (λ=0.999972)

types 7, 8 or 9 and rises to a MELD score of 20 when there is a very low-quality liver such as Organ 10. In general, as the liver quality drops the control limit is nondecreasing and thus the "Wait" region enlarges. This result is intuitive: as the patient's post-transplant life expectancy drops while her pre-transplant life expectancy remains the same, she chooses to wait until she reaches a sicker state.

Disease and patient types affect the optimal policies. Figure 5.5 shows the optimal policy for seven patient types, where patient type is differentiated by characteristics not including the MELD score, ten levels of liver quality and three disease groups, with an annual discount rate of 0.99 as well as 1. Table 5.3 shows the static characteristics of the patient types that are used in generating Figure 5.5. Disease groups in Figure 5.5 are defined in Table A.1. In Figure 5.5, patients are in decreasing order with respect to their post-transplant expected life days for a given liver. As the quality of the patient characteristics drops, the "Wait" region enlarges. Intuitively, this result makes sense because, as the quality of patient characteristics drops, the ratio of pre-transplant survival to post-transplant survival rate increases. In general, our findings indicate that Disease Group 3, which includes the acute liver diseases, has the highest control limits for the same patients and livers.

Table 5.2: Donor Characteristics

Organ No	Age	Sex	Blood Type	Donor white[1]	CMVGR[2]
1	20	Female	A	No	No
2	25	Female	A	No	No
3	35	Female	A	No	No
4	20	Male	A	No	No
5	30	Male	A	No	No
6	40	Male	A	No	No
7	50	Male	A	No	No
8	52	Female	O	No	Yes
9	60	Male	O	No	No
10	70	Male	O	No	Yes

[1] Indicates whether the donor is white or not
[2] Indicates whether the donor has cytomegalovirus (CMVGR) or not

(a)Annual discount rate=0.99 (b)Annual discount rate=1

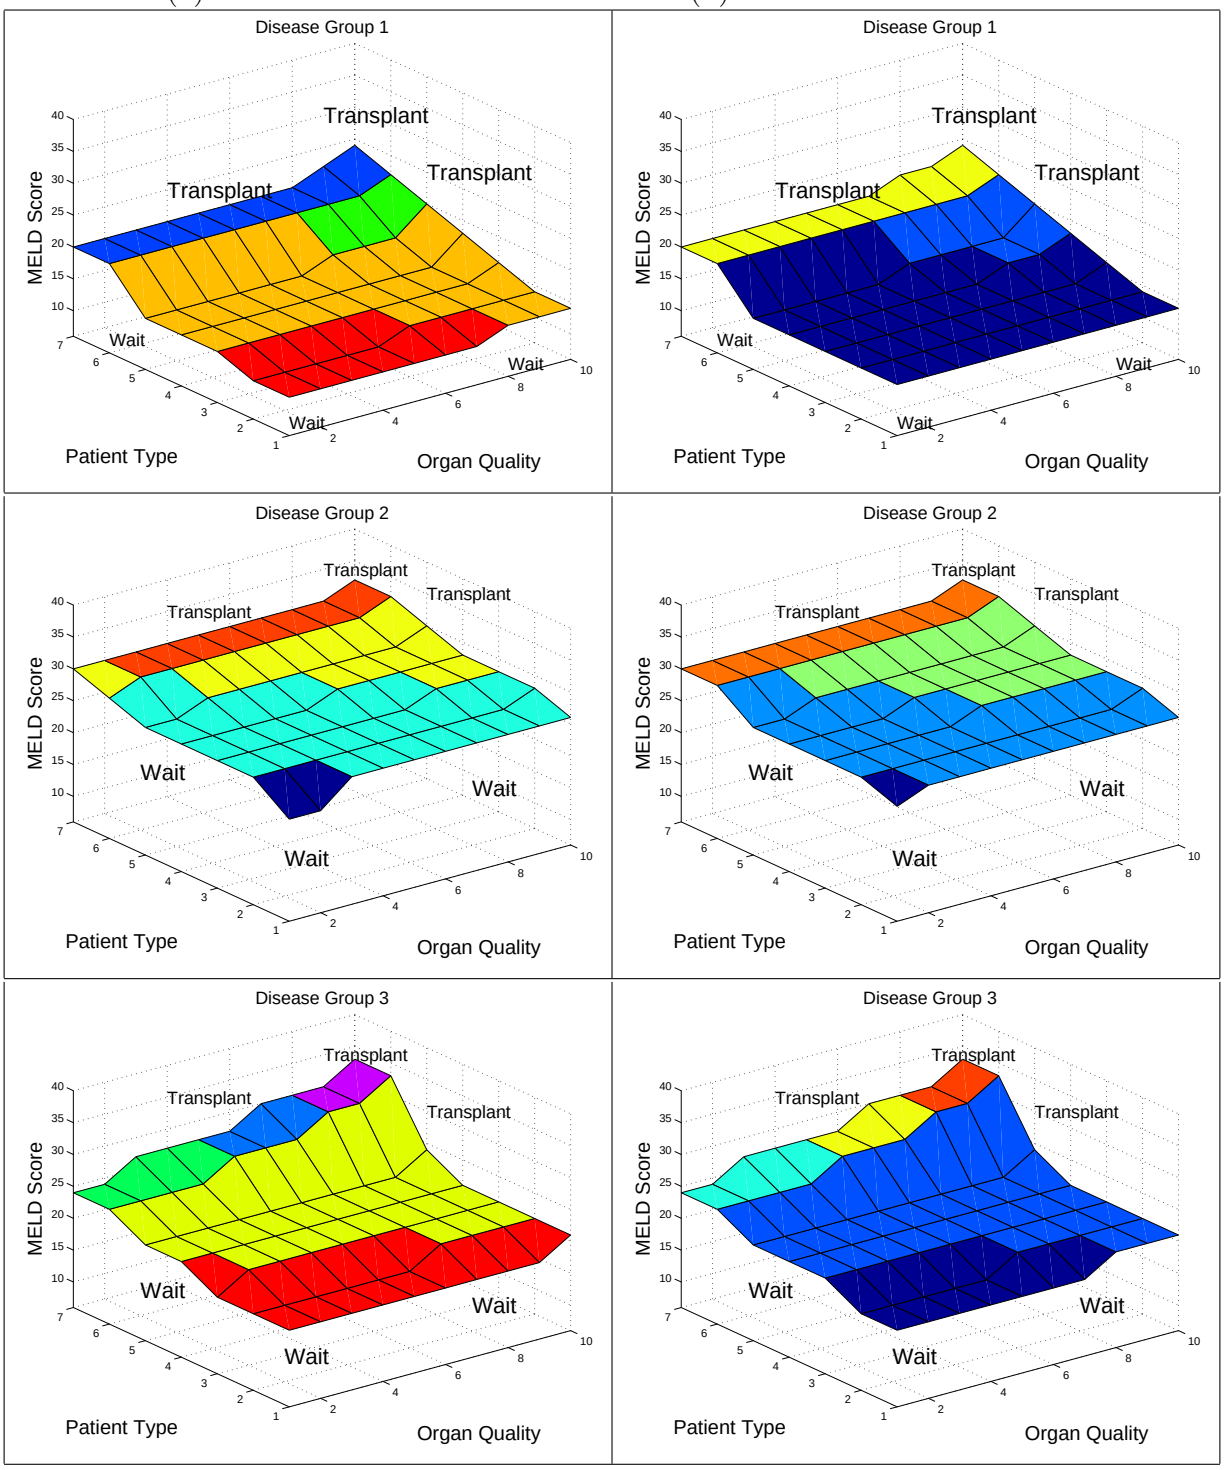

Figure 5.5: Transplant-Wait Decisions for Discounted and Undiscounted Problems

Table 5.3: Patient Characteristics

Patient No[1]	Age	Sex	Blood Type	CMVGR[2]	Encephel[3]	Priortx[4]
1	22	Female	A	No	No	No
2	30	Female	A	No	No	No
3	40	Female	A	No	Yes	No
4	50	Female	A	No	Yes	No
5	60	Female	A	No	Yes	No
6	65	Female	A	Yes	Yes	Yes
7	72	Female	A	Yes	Yes	Yes

[1] The patients are ordered according to their total life expectancies with a given organ, i.e., Patient 1 has the longest expected life and Patient 7 has the shortest life expectancy given the same liver.
[2] Indicates whether the patient has cytomegalovirus (CMVGR) or not
[3] Indicates whether the patient has encephalopathy or not
[4] Indicates whether the patient had a prior transplant or not

We also perform a one-way sensitivity analysis of discount factor for Patient 6. We use different annual discount factors and find the resulting optimal policies. If the only organ available to the Patient 6 is Organ 1 and the patient has hepatitis, the resulting optimal policies for different annual discount rates are presented in Figure 5.6. Note that, Figure 5.6 shows the optimal policies at several discount factors, namely at 0.5, 0.6, 0.7, 0.75, 0.8, 0.85, 0.9, 0.95 and 1. The results of the sensitivity analysis show that the larger the annual discount rate, the larger the "Wait" region. This result is intuitive: as the discount rate increases, the future life days of the patient become more valuable. Note also that the risk of operative mortality is very high because the patient is elderly. Furthermore, the expected post-transplant life days for this particular patient is very limited. In other words, the reward of taking the risk of operation and having transplantation is very small. Hence, the patient might not prefer to take the risk of operation. In addition to these, it is possible that when the patient is relatively healthy, she does not have a high risk of immediate death so, she may not choose to use her option for "Transplant" until she becomes really sick. Note that as long as she waits until she gets sicker, she does not lose much of her post-transplant reward because, future benefit of transplantation is not devalued greatly. On the other hand, if

annual discount rate is low, the present life of the patient will be more valuable than her future life days. As a result, if she does not have the transplantation, the probability of death is immediate; therefore, she might prefer to take the risk of operation which guarantees a certain amount of time for living longer.

5.4 CONCLUSIONS

This section considers the problem of optimally timing living-donor liver transplantation. We formulate an MDP model of this decision problem and solve numerical examples using clinical data. To the best of our knowledge, this study is the first to formulate the living-donor organ transplantation problem by explicitly modeling patient health and the first to solve numerical instances of the problem using clinical data.

We derive structural properties of the LDM, including conditions that guarantee the existence of a control-limit policy. We also establish sufficient conditions for some of the intuitive results seen in our computational experiments. For instance, if one disease causes a faster deterioration in patient health than another, and yet results in identical post-transplant life expectancy, then the control limit for this disease is less than or equal to that for the other.

In all of our computational tests, the optimal policy is of control-limit type. In some of the examples, when the liver quality is very low, it is optimal for the patient to choose never to have the transplant. This implies that measuring an allocation system based on total number of transplants may not fully capture the patient's perspective. Our computational experiments also show that there are significant differences in the optimal policies for identical patients in different disease groups. There are two possible sources of this variation: differences in the disease progression and differences in the post-transplant survival rate of the diseases. We also perform a sensitivity analysis on the discount rate, which shows that as the discount rate increases the number of health states that it is optimal to "Wait" also increases.

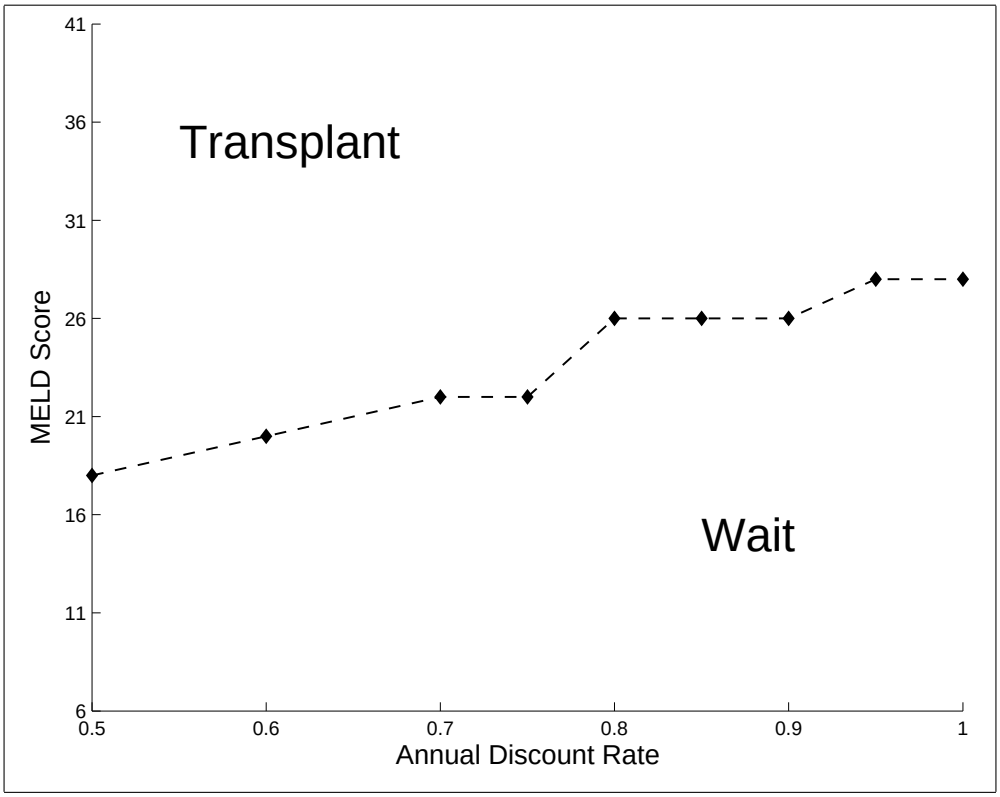

Figure 5.6: Optimal Policies for Different Discount Factors

6.0 CADAVERIC-DONOR-ONLY MODEL

The LDM considers the decision problem when the patient has a living donor who is available at all points in time. Unfortunately, the LDM is not applicable to the cadaveric-donor-only problem, because there are multiple organ types in the cadaveric-donor-only problem, whereas LDM assumes that there is a single organ type available to the patient. Furthermore, cadaveric organ arrivals are random, whereas a living-donor organ is always available. Capturing these complexities requires a different state definition and a different transition probability structure. This chapter describes the Cadaveric-Donor-Only Model (CDM), which formulates and solves the cadaveric-donor-only problem. LDM is a special case of CDM, in which there is a single type of cadaveric liver that is offered to the patient at all decision epochs.

In this chapter, the decision maker must compete with other patients for cadaveric livers. The presence of other patients on the list is reflected in the probability of obtaining offers for livers. If there are many other patients similar to the decision maker, the probability of receiving future offers may be smaller than if there are few similar patients. The outcome of this model is a policy that maps health and liver states to two possible actions: accept a liver offered at a given time, or continue waiting for a liver.

The remainder of this chapter is organized as follows. In Section 6.1, we present an MDP model of the problem. We derive some structural properties of this MDP model and its optimal policy in Section 6.2. In Section 6.3, we present and discuss computational results. We draw some conclusions in Section 6.4.

6.1 MODEL FORMULATION

This problem is the special case of the model presented in Chapter 3, in which $r(h, \ell_{LD}, T_{LD}) = 0, h \in S_H$. That is, the patient does not have a living-donor liver.

Figure 6.1 shows the state-transition diagram of the MDP. The decision maker can take one of two actions in state (h, ℓ), namely, "Accept" the liver ℓ or "Wait for one more decision epoch". If the patient chooses "Accept" in state (h, ℓ), she receives a reward of $r(h, \ell, T_C)$, quits the process and moves to absorbing state "Transplant" with probability 1. If the patient chooses to "Wait" in state (h, ℓ), then she receives an intermediate reward of $r(h, W)$ and moves to state $(h', \ell') \in S$ with probability $\mathcal{P}(h', \ell'|h, \ell)$. The optimal solution to this problem is obtained by solving the following set of recursive equations [125]:

$$V(h,\ell) = \max\left\{ r(h,\ell,T_C), r(h,W) + \lambda \sum_{(h',\ell')\in S} \mathcal{P}(h',\ell'|h,\ell) V(h',\ell') \right\}, \quad h \in \{1,\ldots,H\}, \ell \in S_L. \tag{6.1}$$

6.2 STRUCTURAL PROPERTIES

In this section, we derive some structural properties of the cadaveric-donor-only model (CDM) given by (6.1). The following assumptions are common to all of the theorems:

(As1): The function $r(h, \ell, T_C)$ is nonincreasing in both h and ℓ. That is, as the patient gets sicker and/or the liver quality drops, the patient's post-transplant reward does not increase.

(As2): The function $r(h, W)$ is nonincreasing in h. That is, as the patient gets sicker the intermediate reward does not increase.

Theorem 6.1 proves that the above assumptions guarantee the monotonicity of the optimal value function in liver state for fixed health, without making any additional assumptions on the transition probability matrices.

Theorem 6.1. *$V(h, \ell)$ is monotonically nonincreasing in ℓ, $\ell \in S_L$, $\forall h \in S_H$.*

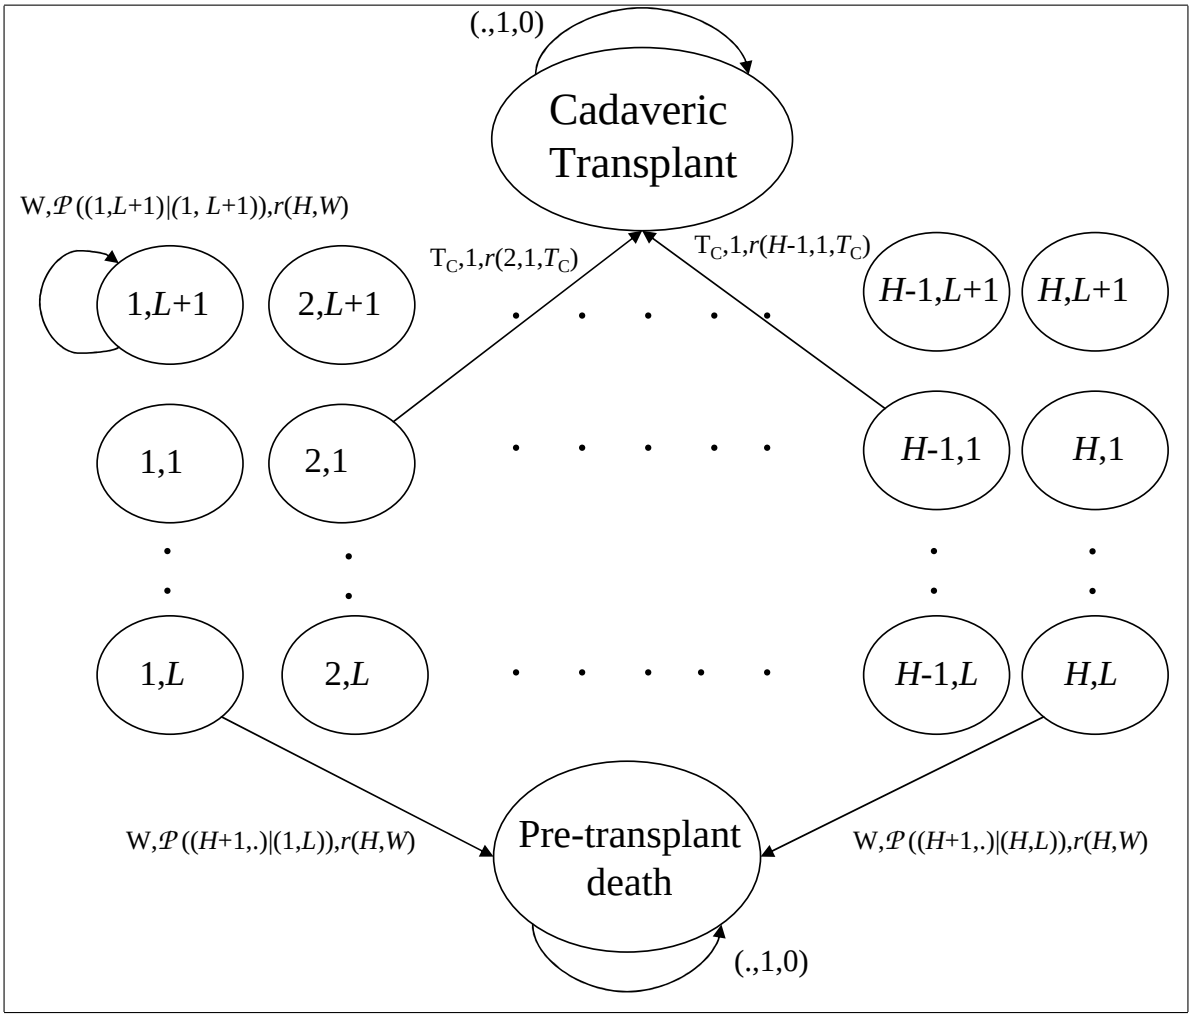

Figure 6.1: State-Transition Diagram of the Cadaveric-Donor-Only Model*

*The labels $a(h,\ell), \mathcal{P}(j|h,\ell), r(h,\ell,a))$ on each arc represent the action taken at state (h,ℓ), the probability that the patient will move to state j when her current state is (h,ℓ), and the reward obtained by taking action a in state (h,ℓ), respectively. Note that this figure does not show all possible transitions.

Proof. If $a^*(h,\ell+1)$='$\mathrm{T_C}$', then $V(h,\ell) \geq r(h,\ell,T_C) \geq r(h,\ell+1,T_C) = V(h,\ell+1)$ from *(As1)* and the result immediately follows. Otherwise,

$$V(h,\ell) \geq r(h,W) + \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h',\ell'|h,\ell) V(h',\ell') \qquad \text{and}$$
$$V(h,\ell+1) = r(h,W) + \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h',\ell'|h,\ell+1) V(h',\ell').$$

Because $\mathcal{P}(h',\ell'|h,\ell+1) = \mathcal{P}(h',\ell'|h,\ell)$, the right hand sides are the same, yielding $V(h,\ell) \geq V(h,\ell+1)$, from which the result follows. □

Although $V(h,\ell)$ is monotonically nonincreasing in ℓ, $V(h,\ell)$ may not be monotonic in h, even if we assume that $\mathcal{H}$ is IFR. Remark 6.1 formalizes this situation.

Remark 6.1. *If $\mathcal{H}$ is IFR, then $V(h,\ell)$ may not be monotonically nonincreasing in h, $h \in S_H$, $\forall \ell \in S_L$.*

Counterexample Let L=1, H=2, λ= 0.99, $r(1,1,T_C)$=$r(2,1,T_C)$=12, $r(1,W) = r(2,W)$=1 and

$$\mathcal{H} = \begin{bmatrix} 0.9 & 0.01 & 0.09 \\ 0 & 0.91 & 0.09 \end{bmatrix} \text{ and } \mathcal{L} = \begin{bmatrix} 0.01 & 0.99 \\ 0.99 & 0.01 \end{bmatrix}.$$

Then the optimal value function is:

$V(1,1)$=12.00 $V(1,2)$=10.40

$V(2,1)$=12.00 $V(2,2)$=11.81

Hence, although $\mathcal{H}$ is IFR and *(As1)* and *(As2)* are satisfied, $V(h,\ell)$ is nonmonotonic in h.

Note that the remark holds if we replace the IFR assumption with the TP_2 assumption. An example for this case is presented in Figure 5.3. Note that the problem in Figure 5.3 can be formulated as a CDM problem if we use the following $\mathcal{L}$ matrix:

$$\mathcal{L} = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}.$$

Therefore the problem in Figure 5.3 constitutes a counterexample for Remark 6.1.

Theorem 6.2 proves the monotonicity of $V(h,\ell)$ in h under a set of intuitive conditions. We first present two lemmas that are used in the proof of Theorem 6.2. Let $V^i(h,\ell)$ be the

value and $a^i(h,\ell)$ be the optimal action associated with state (h,ℓ) at the ith iteration of the value iteration algorithm, respectively. Then, noting that $a^i(h,L+1)$='W', Lemma 6.1 states that at the ith iteration of the value iteration algorithm the optimal value of the liver states for which the optimal action is to "Wait" for the same states is equal to $V^i(h,L+1)$.

Lemma 6.1. *If $a^i(h,\ell)$ ='W' then $V^i(h,\ell) = V^i(h,L+1)$.*

Proof. The proof is obvious and is omitted. □

Lemma 6.2. *Let $z^i(h) = \sum_{\ell\in S_L} \mathcal{L}(\ell|h)V^i(h,\ell)$. If $V^i(h,\ell)$ is nonincreasing in h and ℓ and*

$$\frac{\mathcal{L}(\ell|h+1)}{\mathcal{L}(\ell|h)} \leq \frac{r(h,\ell,T_C)}{r(h+1,\ell,T_C)} \text{ for } h=1,\ldots,H-1 \text{ and } \ell=1,\ldots,L, \tag{6.2}$$

then $z^i(h)$ is nonincreasing in h.

Proof. Note that if $a^i(h,\ell)$='$\mathrm{T_C}$', then $V^i(h,\ell) = r(h,\ell,T_C)$ and if $a^i(h,\ell)$='W', $V^i(h,\ell) = V(h,L+1)$ from Lemma 6.1. Let $L_g^i(h+1) = \{\ell \in S_L : a^i(h+1,\ell) = \text{'}\mathrm{T_C}\text{'}, \mathcal{L}(\ell|h+1) \geq \mathcal{L}(\ell|h)\}$ and $L_s^i(h+1) = \{\ell \in S_L : a^i(h+1,\ell) = \text{'}\mathrm{T_C}\text{'}, \mathcal{L}(\ell|h+1) < \mathcal{L}(\ell|h)\}$. Then,

$$\begin{aligned} z^i(h) - z^i(h+1) \;=\; & \sum_{\ell\in L_g^i(h+1)} [\mathcal{L}(\ell|h)V^i(h,\ell) - \mathcal{L}(\ell|h+1)V^i(h+1,\ell)] + \\ & \sum_{\ell'\in L_s^i(h+1)} [\mathcal{L}(\ell'|h)V^i(h,\ell') - \mathcal{L}(\ell'|h+1)V^i(h+1,\ell')] + \\ & \sum_{\ell''\notin\{L_g^i(h+1)\cup L_s^i(h+1)\}} [\mathcal{L}(\ell''|h)V^i(h,\ell'') - \mathcal{L}(\ell''|h+1)V^i(h+1,\ell'')] \end{aligned} \tag{6.3}$$

Note that $V^i(h+1,\ell) = r^i(h+1,\ell)$ for $\ell \in \{L_g^i(h+1)\cup L_s^i(h+1)\}$ since $a^i(h+1,\ell)$='$\mathrm{T_C}$', and $\mathcal{L}(\ell|h)V^i(h,\ell) - \mathcal{L}(\ell|h+1)V^i(h+1,\ell) \geq 0$ holds for all $\ell \in L_g^(h+1)$ because of (6.2) and *(As1)*. Therefore, the first term on the right hand side of (6.3) can be dropped without increasing the right hand side value. Similarly, $\mathcal{L}(\ell'|h)V^i(h,\ell') - \mathcal{L}(\ell'|h+1)V^i(h+1,\ell') \geq \mathcal{L}(\ell'|h)V^i(h+1,L+1) - \mathcal{L}(\ell'|h+1)V^i(h+1,L+1)$ holds for $\ell' \in L_s^i(h+1)$ because $V^i(h,\ell)$ is monotonic in h and ℓ and $\mathcal{L}(\ell'|h) \geq \mathcal{L}(\ell'|h+1)$ by definition. We can also replace each $V^i(h,\ell'')$ with $V^i(h+1,L+1)$ for $\ell'' \notin \{L_g^i(h+1)\cup L_s^i(h+1)\}$ without changing the value of the right hand side of (6.3) as a result of Lemma 6.1 and the monotonicity of $V(h,l)$ in h and l. Making these changes results in the following:

$$
\begin{aligned}
z^i(h) - z^i(h+1) \;\geq\; & \sum_{\ell' \in L_s^i(h+1)} [\mathcal{L}(\ell'|h) - \mathcal{L}(\ell'|h+1)] V^i(h+1, L+1) + \\
& \sum_{\ell'' \notin \{L_g^i(h+1) \cup L_s^i(h+1)\}} [\mathcal{L}(\ell''|h) V^i(h+1, L+1) - \\
& \mathcal{L}(\ell''|h+1) V^i(h+1, L+1)] \\
= \; & \sum_{\ell' \notin L_g^i(h+1)} [\mathcal{L}(\ell'|h) - \mathcal{L}(\ell'|h+1)] V^i(h+1, L+1) \qquad (6.4)
\end{aligned}
$$

Since $\sum_{\ell' \in L_g(h+1)} \mathcal{L}(\ell'|h) \leq \sum_{\ell' \in L_g(h+1)} \mathcal{L}(\ell'|h+1)$ by definition of $L_g^i(h+1)$, the last inequality is nonnegative from which the result follows. □

Condition (6.2) means that for any given liver type, as the patient gets sicker the increase in the probability of receiving an offer must be smaller than the reduction in the benefit of total expected discounted post-transplant reward. Theorem 6.2 imposes Condition (6.2) on $r(h, \ell, T_C)$ and $\mathcal{L}$ and proves the monotonicity of the optimal value function in h. Using Lemma 6.2, the proof of this theorem is very similar to the proof of Theorem 5.2.

Theorem 6.2. *If $\mathcal{H}$ is IFR and (6.2) holds for all h and ℓ, then $V(h, \ell)$ is nonincreasing in h.*

Proof. We show that at any step of the value iteration algorithm the monotonicity of the value function is preserved. Since it is well known that the value iteration algorithm converges to the optimal value function [125], the result follows.

Assume that the value iteration algorithm starts with a value of 0 for each state, i.e. $V^0(h, \ell) = 0$ for all $(h, \ell) \in S$. For the base case, we need to prove that for a given liver state ℓ, $V^1(h, \ell) \geq V^1(h+1, \ell)$, for $h = 1, \ldots, H$, where

$$
V^1(h, \ell) = \max \left\{ r(h, \ell, T_C), r(h, W) + \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h', \ell'|h, \ell) V^0(h', \ell') \right\}.
$$

Because $V^0(h', \ell') = 0$ for $(h', \ell') \in S$, $V^1(h, \ell) = \max\{r(h, \ell, T_C), r(h, W)\}$. Because both $r(h, \ell, T_C)$ and $r(h, W)$ are monotonically nonincreasing, it is obvious that the value function is also monotonically nonincreasing.

Now, suppose that the property holds for iterations $i = 2, \ldots, n$, that is, $V^i(h, \ell) \geq V^i(h+1, \ell)$ for $h = 1, \ldots, H$ and $i = 2, \ldots, n$. Note that

$$V^{n+1}(h, \ell) = \max \left\{ r(h, \ell, T_C), r(h, W) + \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h', \ell' | h, \ell) V^n(h', \ell') \right\} \quad \text{and}$$

$$V^{n+1}(h+1, \ell) = \max \left\{ r(h+1, \ell, T_C), r(h+1, W) + \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h', \ell' | h+1, \ell) V^n(h', \ell') \right\}.$$

If $V^{n+1}(h+1, \ell) = r(h+1, \ell, T_C)$, $V^{n+1}(h, \ell) \geq r(h, \ell, T_C) \geq r(h+1, \ell, T_C) = V^{n+1}(h+1, \ell)$ and the result immediately follows. Otherwise, we obtain the following:

$$\begin{aligned} & V^{n+1}(h, \ell) - V^{n+1}(h+1, \ell) \\ \geq \; & r(h, W) - r(h+1, W) + \lambda \left[\sum_{h' \leq h} \mathcal{H}(h'|h) z^n(h') + \sum_{h'' > h} \mathcal{H}(h''|h) z^n(h'') \right] - \\ & \lambda \left[\sum_{h' \leq h} \mathcal{H}(h'|h+1) z^n(h') + \sum_{h'' > h} \mathcal{H}(h''|h+1) z^n(h'') \right] \\ \geq \; & \lambda \left(\sum_{h' \leq h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)] z^n(h') + \sum_{h'' > h} [\mathcal{H}(h''|h) - \mathcal{H}(h''|h+1)] z^n(h'') \right), \quad (6.5) \end{aligned}$$

where (6.5) follows because of *(As2)*. From the induction assumption and Lemma 6.2, $z^n(h)$ is monotonic. Therefore, as a result of Lemma 5.1, the inequality is not violated when we replace $z^n(h')$ with $z^n(h)$ for $h' \leq h$ and $z^n(h'')$ with $z^n(h+1)$ for $h'' > h$. Hence, the above inequality can be rewritten as follows:

$$\begin{aligned} & V^{n+1}(h, \ell) - V^{n+1}(h+1, \ell) \\ \geq \; & \lambda \left(\sum_{h' \leq h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)] z^n(h) + \sum_{h'' > h} [\mathcal{H}(h''|h) - \mathcal{H}(h''|h+1)] z^n(h+1) \right) \\ = \; & \lambda \left(\sum_{h' \leq h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)] [z^n(h) - z^n(h+1)] \right). \end{aligned}$$

By the IFR property, the induction assumption and Lemma 6.2, both $z^n(h) - z^n(h+1)$ and $\sum_{h' \leq h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)]$ are nonnegative. Therefore, $V^{n+1}(h) - V^{n+1}(h+1)$ is also nonnegative, from which the result follows. □

Unlike the case of the LDM, there are two types of control limits in the CDM: liver-based and health-based.

Definition 6.1. *A* liver-based control-limit policy *is of the following form: For a given health state h, choose the "Transplant" action and "Accept" the liver if and only if the offered liver is of type $1, 2, \ldots, i(h)$, for some liver state $i(h)$, called the* liver-based control limit.

Similarly, a health-based control-limit policy *is of the simple form: For a given liver state ℓ, choose the "Transplant" action and "Accept" the liver if and only if the health state is one of the states $j(\ell), j(\ell)+1, \ldots, H$, for some health state $j(\ell)$, called the* health-based control limit.

Theorem 6.3. *There exists an optimal liver-based control-limit policy.*

Proof. The theorem is equivalent to showing that $a^*(h, \ell+1)$='$\mathrm{T_C}$' implies $a^*(h,\ell)$='$\mathrm{T_C}$', $\forall h \in S_H$. If $a^*(h,\ell+1)$='$\mathrm{T_C}$' for a given $h \in S_H$, then the following are true about the optimality equations:

$$V(h,\ell+1) = r(h,\ell+1,T_C) \geq r(h,W) + \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h',\ell'|h,\ell+1) V(h',\ell'),$$

and

$$V(h,\ell) = \max \left\{ r(h,\ell,T_C), r(h,W) + \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h',\ell'|h,\ell) V(h',\ell') \right\}.$$

It is obvious that

$\lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h',\ell'|h,\ell+1) V(h',\ell') = \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h',\ell'|h,\ell) V(h',\ell')$ holds in the above expressions. As a result of *(As1)* we can write the following:

$$r(h,\ell,T_C) \geq r(h,\ell+1,T_C) \geq r(h,W) + \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h',\ell'|h,\ell) V(h',\ell'),$$

which is equivalent to $a^*(h,\ell)$='$\mathrm{T_C}$', from which the result follows. □

While *(As1)* and *(As2)* suffice to guarantee the existence of an optimal liver-based control limit policy, we need additional assumptions to prove the existence of an optimal health-based liver policy. The following theorem gives a set of intuitive conditions that ensures the existence of an optimal health-based control-limit policy. Note that the conditions of this theorem and those of Theorem 5.3 have similar interpretations.

Theorem 6.4. *If $\mathcal{H}$ is IFR, $\mathcal{L}$ satisfies (6.2),*

$$\sum_{k=j}^{H} \mathcal{H}(k|h) \leq \sum_{k=j}^{H} \mathcal{H}(k|h+1), \text{ for } j = h+1, \ldots, H \text{ and } h = 1, \ldots, H, \tag{6.6}$$

$$\text{and } \frac{r(h,\ell,T_C) - r(h+1,\ell,T_C)}{r(h+1,\ell,T_C)} \leq \lambda[\mathcal{H}(H+1|h+1) - \mathcal{H}(H+1|h)], \text{ for } h = 1, \ldots, H-1, \tag{6.7}$$

then there exists an optimal health-based control-limit policy.

Proof. First note that the monotonicity results of Theorems 6.1 and 6.2 hold. If we show that for some h, $a^*(h,\ell)$='$\mathrm{T_C}$' implies $a^*(h+1,\ell)$='$\mathrm{T_C}$' then the result follows. Assume that the converse is true. In other words, for some h, $a^*(h,\ell)$='$\mathrm{T_C}$' but $a^*(h+1,\ell)$ is uniquely 'W'. In this case, $r(h,\ell,T_C) \geq r(h,W) + \lambda \sum_{h' \in S_H} \mathcal{H}(h'|h) z(h')$ and $r(h+1,\ell,T_C) < r(h+1,W) + \lambda \sum_{h' \in S_H} \mathcal{H}(h'|h+1) z(h')$. It follows that,

$$\begin{aligned}
r(h,\ell,T_C) - r(h+1,\ell,T_C) \;>\; & r(h,W) - r(h+1,W) + \\
& \lambda \left[\sum_{h' \leq h} \mathcal{H}(h'|h) z(h') + \sum_{h'' > h} \mathcal{H}(h''|h) z(h'') \right] - \\
& \lambda \left[\sum_{h' \leq h} \mathcal{H}(h'|h+1) z(h') + \sum_{h'' > h} \mathcal{H}(h''|h+1) z(h'') \right], \\
\geq \; & \lambda \left[\sum_{h' \leq h} \mathcal{H}(h'|h) z(h') + \sum_{h''=h+1}^{H} \mathcal{H}(h''|h) z(h'') \right] - \\
& \lambda \left[\sum_{h' \leq h} \mathcal{H}(h'|h+1) z(h') + \sum_{h''=h+1}^{H} \mathcal{H}(h''|h+1) z(h'') \right],
\end{aligned}$$

because of *(As2)* and $z(H+1) = 0$. The last inequality can also be rewritten as

$$\begin{aligned}
r(h,\ell,T_C) - r(h+1,\ell,T_C) > \lambda \Bigg(& \sum_{h' \leq h} [\mathcal{H}(h'|h) - \mathcal{H}(h'|h+1)] z(h') + \\
& \sum_{h''=h+1}^{H} [\mathcal{H}(h''|h) - \mathcal{H}(h''|h+1)] z(h'') \Bigg).
\end{aligned} \tag{6.8}$$

From Lemma 6.2 we know that $z(h)$ is nonincreasing in h. Therefore, the result of Lemma 5.1 applies to this problem and we can replace each $z(h')$ with $z(h)$ without violating (6.8).

Similarly, as a result of Lemma 5.1 and (6.6), we can replace each $z(h'')$ with $z(h+1)$. We then obtain the following:

$$\begin{aligned} r(h,\ell,T_C) - r(h+1,\ell,T_C) &> \lambda \Bigg(\Bigg[\Bigg(1 - \sum_{h''=h+1}^{H} \mathcal{H}(h''|h) - \mathcal{H}(H+1|h) \Bigg) - \\ &\quad \Bigg(1 - \sum_{h''=h+1}^{H} \mathcal{H}(h''|h+1) - \mathcal{H}(H+1|h+1) \Bigg) \Bigg] z(h) + \\ &\quad \sum_{h''=h+1}^{H} [\mathcal{H}(h''|h) - \mathcal{H}(h''|h+1)] z(h+1) \Bigg) \qquad (6.9) \\ &= \lambda \Bigg(\sum_{h''=h+1}^{H} [\mathcal{H}(h''|h+1) - \mathcal{H}(h''|h)][z(h) - z(h+1)] \Bigg) + \\ &\quad \lambda \left([\mathcal{H}(H+1|h+1) - \mathcal{H}(H+1|h)] z(h) \right). \end{aligned}$$

From (6.6), $\sum_{h''=h+1}^{H} [\mathcal{H}(h''|h+1) - \mathcal{H}(h''|h)]$ is nonnegative, so we can drop the first term and rewrite the last inequality as follows:

$$r(h,\ell,T_C) - r(h+1,\ell,T_C) > \lambda [\mathcal{H}(H+1|h+1) - \mathcal{H}(H+1|h)] z(h). \qquad (6.10)$$

Using (6.7) and (6.10) we obtain $\lambda[\mathcal{H}(H+1|h+1) - \mathcal{H}(H+1|h)]z(h) < \lambda[\mathcal{H}(H+1|h+1) - \mathcal{H}(H+1|h)]r(h+1,\ell,T_C)$, which is equivalent to assuming that $z(h) < r(h+1,\ell,T_C)$. However, Theorem 6.1 and $a^*(h+1,\ell) =$'W' imply that $r(h+1,\ell) < V(h+1,\ell) = V(h+1,L+1) \leq V(h,L+1) \leq z(h)$, because $V(h,L+1) \leq V(h,\ell)$ for $\ell \in S_L$, which means there exists a contradiction. Therefore, $a^*(h+1,\ell)$='$\mathrm{T_C}$' also holds, from which the result follows. □

Theorem 6.5 compares the optimal control limits of two similar patients who are listed in different organ procurement organizations (OPO)s, i.e. who have different liver offer probabilities. Namely, if a patient receives more and better liver offers than another patient, then her optimal health-based control limits will be higher than the second patient given that they have identical rewards and health transition probabilities. Similarly, the liver-based optimal control limits of the patient who receives more and better liver offers will be lower than those of the second one. Intuitively, this is because the patient receiving more and better liver offers will be more selective than the second one.

Theorem 6.5. *Let Π_1 and Π_2 be two instances with liver transition probability matrices $\mathcal{L}_1$ and $\mathcal{L}_2$, respectively. Let V_1 and V_2 be the optimal value functions of Π_1 and Π_2, respectively. If Π_1 and Π_2 have the same reward functions, $r(h, \ell, T_C)$ and $r(h, W)$, the same health transition probability matrix $\mathcal{H}$ and $\mathcal{L}_1 \succeq \mathcal{L}_2$, then the following are true:*

(a) $V_1(h, \ell) \geq V_2(h, \ell)$ for $h \in S_H$ and $\ell \in S_L$.

(b) Let $i_1(h)$ and $i_2(h)$, $h \in S_H$ be the liver-based control limits of Π_1 and Π_2, respectively. Then $i_1(h) \leq i_2(h)$, for all $h \in S_H$.

(c) If both Π_1 and Π_2 have health-based control-limit optimal policies with health-based control limits $j_1(\ell)$ and $j_2(\ell)$, $\ell \in S_L$, then $j_1(\ell) \geq j_2(\ell)$, for all $\ell \in S_L$.

Proof. (a) Suppose that we solve the two problems simultaneously using the value iteration algorithm. We first show that starting with a value of 0 for all states in both problems, at the end of each iteration of the algorithm, the value function of Π_1 will be greater than or equal to the value function of Π_2 for each state. Let $V_i^j(h, \ell)$ be the value function of the state (h, ℓ) of problem i at the end of iteration j. We start with 0 for both problems. In this case, $V_1^1(h, \ell) = V_2^1(h, \ell) = \max\{r(h, \ell, T_C), r(h, W)\}$, for $h \in S_H$ and $\ell \in S_L$ as shown in Theorem 6.2. Since the maximum of two nonincreasing functions is also a nonincreasing function, the result holds for the base case.

Now, assume that $V_1^n(h, \ell) \geq V_2^n(h, \ell)$, $(h, \ell) \in S$ holds for iterations $2, \ldots, n$. Then we want to show that $V_1^{n+1}(h, \ell) \geq V_2^{n+1}(h, \ell)$, $(h, \ell) \in S$. If for any state $(h, \ell) \in S$, $V_2^{n+1}(h) = r(h, \ell, T_C)$, then the result immediately follows since $V_1^{n+1}(h, \ell) \geq r(h, \ell, T_C)$. Otherwise, the application of the value iteration algorithm results in the following:

$$V_1^{n+1}(h, \ell) \geq r(h, W) + \lambda \sum_{h' \in S_H} \mathcal{H}(h'|h) \left[\sum_{\ell' \leq \ell} \mathcal{L}_1(\ell'|h') V_1^n(h', \ell') + \sum_{\ell'' > \ell} \mathcal{L}_1(\ell''|h') V_1^n(h', \ell'') \right]$$

and

$$V_2^{n+1}(h, \ell) = r(h, W) + \lambda \sum_{h' \in S_H} \mathcal{H}(h'|h) \left[\sum_{\ell' \leq \ell} \mathcal{L}_2(\ell'|h') V_2^n(h', \ell') + \sum_{\ell'' > \ell} \mathcal{L}_2(\ell''|h') V_2^n(h', \ell'') \right].$$

We obtain the following:

$$
\begin{aligned}
V_1^{n+1}(h,\ell) - V_2^{n+1}(h,\ell) \;\geq\; & \lambda \sum_{h' \in S_H} \mathcal{H}(h'|h) \Bigg(\sum_{\ell' \leq \ell} \mathcal{L}_1(\ell'|h') V_1^n(h',\ell') + \\
& \sum_{\ell'' > \ell} \mathcal{L}_1(\ell''|h') V_1^n(h',\ell'') - \sum_{\ell' \leq \ell} \mathcal{L}_2(\ell'|h') V_2^n(h',\ell') - \\
& \sum_{\ell'' > \ell} \mathcal{L}_2(\ell''|h') V_2^n(h',\ell'') \Bigg) \\
\geq\; & \lambda \sum_{h' \in S_H} \mathcal{H}(h'|h) \Bigg(\sum_{\ell' \leq \ell} \mathcal{L}_1(\ell'|h') V_2^n(h',\ell') + \\
& \sum_{\ell'' > \ell} \mathcal{L}_1(\ell''|h') V_2^n(h',\ell'') - \sum_{\ell' \leq \ell} \mathcal{L}_2(\ell'|h') V_2^n(h',\ell') - \\
& \sum_{\ell'' > \ell} \mathcal{L}_2(\ell''|h') V_2^n(h',\ell'') \Bigg) \qquad (6.11) \\
=\; & \lambda \sum_{h' \in S_H} \Bigg(\sum_{\ell' \leq \ell} [\mathcal{L}_1(\ell'|h') - \mathcal{L}_2(\ell'|h')] V_2^n(h',\ell') + \\
& \sum_{\ell'' > \ell} [\mathcal{L}_1(\ell''|h') - \mathcal{L}_2(\ell''|h')] V_2^n(h',\ell'') \Bigg) \qquad (6.12) \\
\geq\; & \lambda \sum_{h' \in S_H} \Bigg(\sum_{\ell' \leq \ell} [\mathcal{L}_1(\ell'|h') - \mathcal{L}_2(\ell'|h')] V_2^n(h',\ell) + \\
& \sum_{\ell'' > \ell} [\mathcal{L}_1(\ell''|h') - \mathcal{L}_2(\ell''|h')] V_2^n(h',\ell+1) \Bigg) \qquad (6.13) \\
=\; & \lambda \sum_{h' \in S_H} \Bigg(\sum_{\ell' \leq \ell} [\mathcal{L}_1(\ell'|h') - \mathcal{L}_2(\ell'|h')] [V_2^n(h',\ell) - V_2^n(h',\ell+1)] \Bigg) \qquad (6.14) \\
\geq\; & 0 \qquad (6.15)
\end{aligned}
$$

where (6.11) follows from the induction assumption and (6.12) is obtained by simply rearranging terms. Inequality (6.13) holds because $\mathcal{L}_1 \succeq \mathcal{L}_2$ and the monotonicity of the value function in ℓ imply that each $V_2^n(h',\ell')$ can be replaced with $V_2^n(h',\ell)$, and $V_2^n(h',\ell'')$ can be replaced with $V_2^n(h',\ell+1)$, without violating the inequality as a result of Lemma 5.2. (6.14) follows from rearranging the terms in (6.13); and (6.15) follows because $\mathcal{L}_1 \succeq \mathcal{L}_2$ and the value function is monotonic in ℓ.

Since the value function for Π_1 is always greater than or equal to that of Π_2 at each iteration of the value iteration algorithm, the optimal value function of Π_1 will always be greater than or equal to that of Π_2 from which the result follows.

(b) Note that the value function is monotonic from Part (a). Therefore, for a given h, if for state $i_1(h)$, $a_1(h, i_1(h))$='$\mathrm{T_C}$' in the first problem, since $r(h, i_1(h), T_C) \leq V_2(h, i_1(h)) \leq V_1(h, i_1(h)) = r(h, i_1(h), T_C)$, $a_2(h, i_1(h))$='$\mathrm{T_C}$' always holds in the second problem. As a result, since we have a liver-based control-limit optimal policy in both problems, the liver-based control limit in the first problem will always be smaller than or equal to the liver-based control limit in the second problem, from which Part (b) of the theorem follows.

(c) Similarly, for a given ℓ, if for state $j_1(\ell)$, $a_1(j_1(\ell), \ell)$='$\mathrm{T_C}$' in the first problem, since $r(j_1(\ell), \ell, T_C) \leq V_2(j_1(\ell), \ell) \leq V_1(j_1(\ell), \ell) = r(j_1(\ell), \ell, T_C)$, $a_2(j_1(\ell), \ell)$='$\mathrm{T_C}$' always holds in the second problem. As a result, since we have a health-based control-limit optimal policy in both problems, the health-based control limit in the first problem will always be greater than or equal to the health-based control limit in the second problem, from which the Part (c) of the theorem follows. □

Theorem 6.6 provides a similar result to Theorem 6.5, in which two patients having identical reward functions and liver offer probabilities have different health transition probability matrices. If a patient's health deteriorates faster than another patient, then she will have lower health-based and higher liver-based control limits than the second patient.

Theorem 6.6. *Let Π_1 and Π_2 be two instances with health transition matrices $\mathcal{H}_1$ and $\mathcal{H}_2$, respectively. Let V_1 and V_2 be the optimal value functions of Π_1 and Π_2, respectively. If Π_1 and Π_2 has the same reward functions, $r(h, \ell, T_C)$ and $r(h, W)$, the same liver transition probability matrix $\mathcal{L}$ and $\mathcal{H}_1 \succeq \mathcal{H}_2$, then the following are true:*

(a) $V_1(h, \ell) \geq V_2(h, \ell)$ for $h \in S_H$ and $\ell \in S_L$.

(b) Let $i_1(h)$ and $i_2(h)$, $h \in S_H$ be the liver-based control limits of Π_1 and Π_2, respectively. Then $i_1(h) \leq i_2(h)$, for all $h \in S_H$.

(c) If both Π_1 and Π_2 have heath-based control-limit optimal policies with health-based control limits $j_1(\ell)$ and $j_2(\ell)$, $\ell \in S_L$, then $j_1(\ell) \geq j_2(\ell)$, for all $\ell \in S_L$.

Proof. (a) Suppose that we solve the two problems simultaneously using the value iteration algorithm. We first show that starting with a value of 0 for all states in both problems, at the end of each iteration of the algorithm, the value function of Π_1 will be greater than or equal to the value function of Π_2 for each state. Let $V_i^j(h,\ell)$ be the value function of the state (h,ℓ) of problem i at the end of iteration j. We start with 0 for both problems. In this case, $V_1^1(h,\ell) = V_2^1(h,\ell) = \max\{r(h,\ell,T_C), r(h,W)\}$, for $h \in S_H$ and $\ell \in S_L$ as shown in Theorem 6.2. Since the maximum of two nonincreasing functions is also a nonincreasing function, the result holds for the base case.

Now, assume that $V_1^n(h,\ell) \geq V_2^n(h,\ell),\ (h,\ell) \in S$ holds for iterations $2,\ldots,n$. Then we want to show that $V_1^{n+1}(h,\ell) \geq V_2^{n+1}(h,\ell),\ (h,\ell) \in S$. If for any state $(h,\ell) \in S$, $V_2^{n+1}(h) = r(h,\ell,T_C)$, then the result immediately follows since $V_1^{n+1}(h,\ell) \geq r(h,\ell,T_C)$. Otherwise, the application of the value iteration algorithm results in the following:

$$V_1^{n+1}(h,\ell) \geq r(h,W) + \lambda \left[\sum_{h' \leq h} \mathcal{L}_1(h'|h) z_1^n(h') + \sum_{h'' > h} \mathcal{L}_1(h''|h) z_1^n(h'') \right]$$

and

$$V_2^{n+1}(h,\ell) = r(h,W) + \lambda \left[\sum_{h' \leq h} \mathcal{L}_2(h'|h) z_2^n(h') + \sum_{h'' > h} \mathcal{L}_2(h''|h) z_2^n(h'') \right],$$

where $z_1^n(h) = \sum_{\ell \in S_L} \mathcal{L}(\ell|h) V_1^n(h,\ell)$ and $z_2^n(h) = \sum_{\ell \in S_L} \mathcal{L}(\ell|h) V_2^n(h,\ell)$.

We obtain the following:

$$\begin{aligned}
V_1^{n+1}(h,\ell) - V_2^{n+1}(h,\ell) &\geq \lambda \sum_{h' \leq h} \mathcal{H}_1(h'|h) z_1^n(h') + \lambda \sum_{h'' > h} \mathcal{H}_1(h''|h) z_1^n(h'') - \\
&\quad \lambda \sum_{h' \leq h} \mathcal{H}_2(h'|h) z_2^n(h') - \lambda \sum_{h'' > h} \mathcal{H}_2(h''|h) z_2^n(h'') \\
&\geq \lambda \sum_{h' \leq h} \mathcal{H}_1(h'|h) z_2^n(h') + \lambda \sum_{h'' > h} \mathcal{H}_1(h''|h) z_2^n(h'') - \\
&\quad \lambda \sum_{h' \leq h} \mathcal{H}_2(h'|h) z_2^n(h') - \lambda \sum_{h'' > h} \mathcal{H}_2(h''|h) z_2^n(h'') \\
&= \lambda \left(\sum_{h' \leq h} [\mathcal{H}_1(h'|h) - \mathcal{H}_2(h'|h)] z_2^n(h') + \right. \\
&\quad \left. \sum_{h'' > h} [\mathcal{H}_1(h''|h) - \mathcal{H}_2(h''|h)] z_2^n(h'') \right) \qquad (6.16)
\end{aligned}$$

$$
\begin{aligned}
&\geq \lambda \left(\sum_{h' \leq h} [\mathcal{H}_1(h'|h) - \mathcal{H}_2(h'|h)] z_2^n(h) + \right. \\
&\quad \left. \sum_{h'' > h} [\mathcal{H}_1(h''|h) - \mathcal{H}_2(h''|h)] z_2^n(h+1) \right) && (6.17) \\
&= \lambda \left(\sum_{h' \leq h} [\mathcal{H}_1(h'|h) - \mathcal{H}_2(h'|h)][z_2^n(h) - z_2^n(h+1)] \right) && (6.18) \\
&\geq 0 && (6.19)
\end{aligned}
$$

where (6.16) follows from the induction assumption and (6.16) is obtained by simply re-arranging terms. Inequality (6.17) holds because $\mathcal{H}_1 \succeq \mathcal{H}_2$ and the monotonicity of the z-function imply that $z_2^n(h')$ can be replaced with $z_2^n(h)$, and $z_2^n(h'')$ can be replaced with $z_2^n(h+1)$, without violating the inequality as a result of Lemma 5.2. Inequality (6.18) follows from rearranging the terms in (6.17) and (6.19) follows because $\mathcal{H}_1 \succeq \mathcal{H}_2$ and the z-function is monotonic.

Since the value function for Π_1 is always greater than or equal to that of Π_2 at each iteration of the value iteration algorithm, the optimal value function of Π_1 will always be greater than or equal to that of Π_2 from which the result follows.

(b) The proof is similar to the proof of Part (b) of Theorem 6.5 and is omitted.

(c) The proof is similar to the proof of Part (c) of Theorem 6.5 and is omitted. □

6.3 COMPUTATIONAL RESULTS

We solve the CDM using real data. The estimation of parameters is described in Chapter 4 and Section 5.3. We first consider the same patient considered in Chapter 5, a 60-year old female patient with primary biliary cirrhosis who has blood type A. We apply a 0.99 annual discount rate (daily $\lambda = 0.999972$) and use the national liver probability matrix. Recall that as the MELD score increases, the patient gets sicker. In Figure 6.2, the livers are ordered from highest to lowest quality, i.e. Liver Type 1 is the "best" and Liver Type 14 is the "worst" liver. The quality of a liver for a given patient is determined by the Post-transplant Survival Model described in Section 4.3. As Figure 6.2 shows, the optimal action varies for different liver types. If the patient has a MELD score of 16, then the optimal policy for this

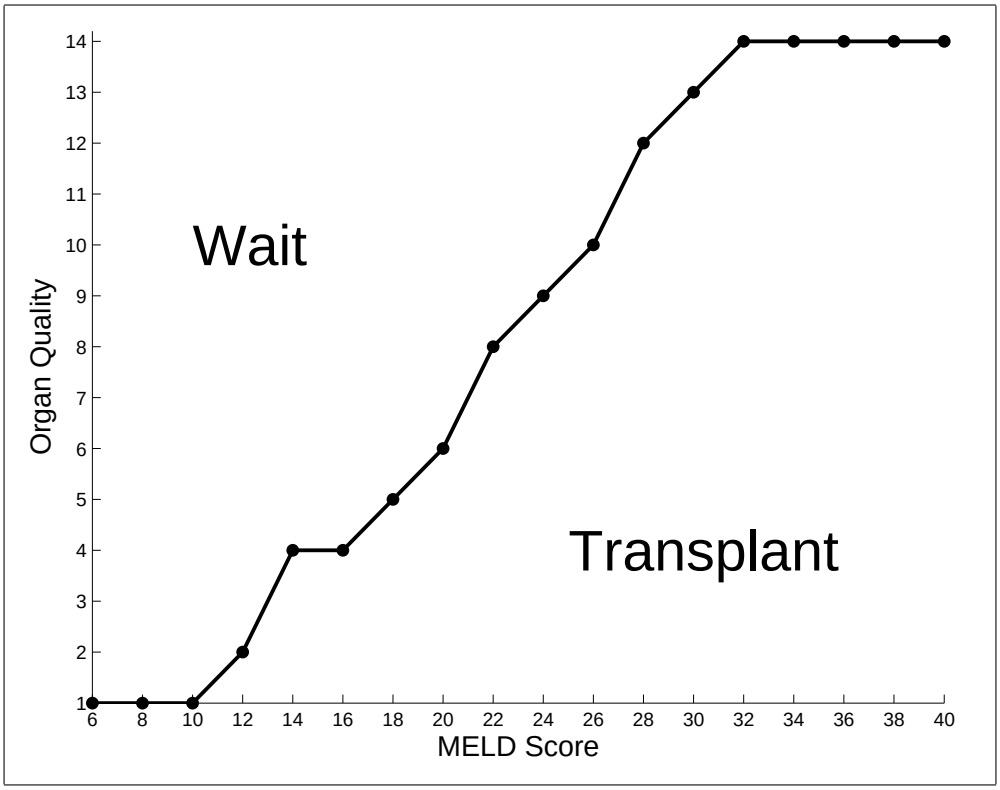

Figure 6.2: Transplant-Wait Decisions for Annual Discount Rate=0.99 (λ=0.999972)

particular patient is as follows: if she receives a liver offer that is of type 4 or worse, then "Wait", otherwise accept the liver and have the transplantation. This policy is an example of a liver-based control-limit policy. Similarly, if the patient receives a liver offer of type 5, then the optimal action for this particular patient is as follows: if her MELD score is below 18, then "Wait", otherwise accept the liver and have the transplantation. This policy is an example of a health-based control-limit policy. For this particular patient, the optimal policy is of liver-based control-limit type for each h and is of health-based control-limit type for each ℓ.

In all of our computational tests, all optimal policies are of liver-based control-limit type. All of our data sets satisfy the requirements of Theorem 6.3. However, some optimal policies are not of health-based control-limit type. Figure 6.3 shows an example of this situation. Figure 6.3 shows the optimal policy for the same patient, but assumes that the patient is listed in region 2 and uses the corresponding liver probability matrix. As can be seen from the figure, for liver type 13, the patient does not have a health-based control-limit policy. This result can be explained intuitively as follows: the current liver allocation system implies that the sicker the patient is, the more likely it is that she receives more frequent and higher quality liver offers. As a result, if as the patient gets sicker the rate of increase in the liver offer probability is very high, then it may be optimal for the patient to decline the low quality livers in anticipation of higher quality liver offers.

Neither the problem in Figure 6.2 nor the problem in Figure 6.3 strictly satisfy the conditions of Theorem 6.4. While the maximum violation for Condition 6.2 is different, the maximum violations for all remaining conditions of Theorem 6.4 are the same for both problems. To quantify the magnitude of the violation of Condition 6.2 we define the following metric:
$\epsilon_4 = \max_{h,\ell} \left\{0, \frac{r(h,\ell,T_C)}{r(h+1,\ell,T_C)} - \frac{\mathcal{L}(\ell|h+1)}{\mathcal{L}(\ell|h)}\right\}$, the maximum violation of Condition 6.2. The value for ϵ_4 is 0.41 and 1.01 for the problems in Figure 6.2 and Figure 6.3, respectively. Although both of these values are high, the problem in Figure 6.3 has a larger ϵ_4 value, which shows that the sufficiency conditions in Theorem 6.4 may provide a good measure for the existence of a health-based control-limit policy.

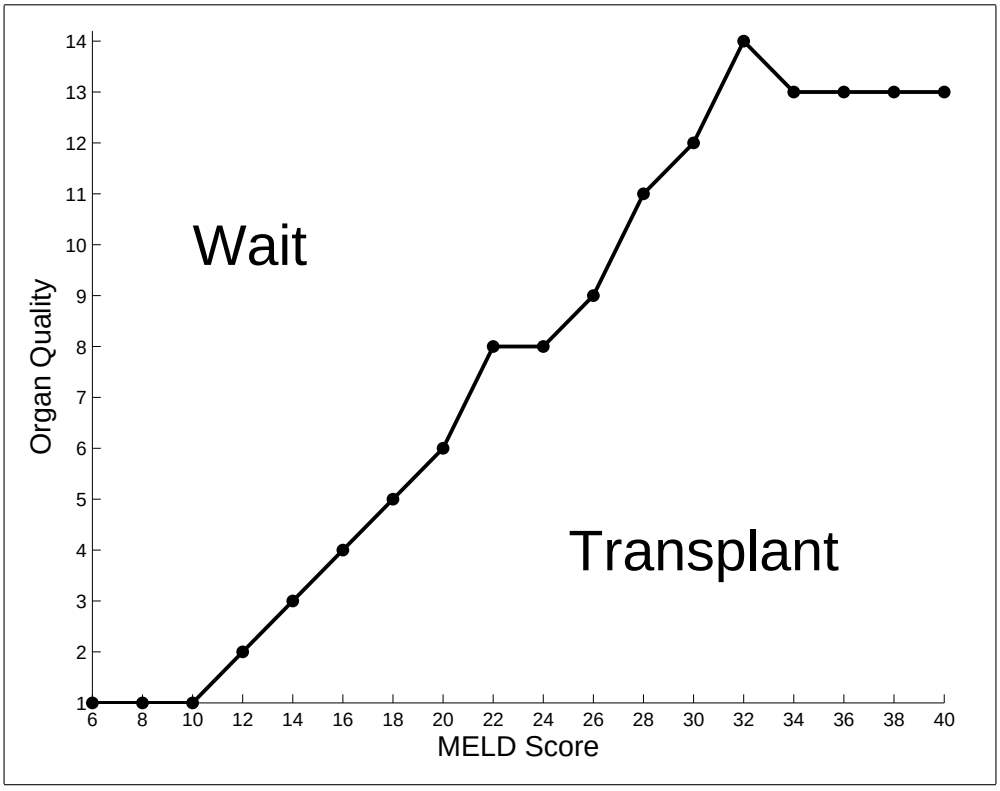

Figure 6.3: Example of a Non-health Based Control-limit Policy for Liver Type 13

As in the LDM, disease and patient types affect the optimal policies. Figure 6.4 shows the optimal policy for seven patient types with an annual discount rate of 0.99 as well as 1. We use the patient type definitions in Table 5.3 to generate Figure 6.4. Disease groups in Figure 6.4 are defined in Table A.1. In Figure 6.4, patients are in decreasing order with respect to their post-transplant expected life days for any given liver. As in the LDM, as the quality of the patient characteristics drops, it is optimal to "Wait" in more states. Intuitively, this result makes sense, because as the quality of patient characteristics drops, the ratio of pre-transplant survival probability to post-transplant survival probability increases. In general, our findings indicate that Disease Group 3, which includes the acute liver diseases, has the highest control limits. Our results also show that as in the LDM, the optimal control limits of the undiscounted problem are higher than those of the discounted problem.

Liver offer probabilities also affect the optimal policy. Recall that the liver offer probability matrix depends on the location where the patient is listed. Figure 6.5 compares the optimal transplant-wait decisions for Patient 5 when she is listed in two different regions. We assume that the patient has cirrhosis and the annual discount rate is 0.99. There is not a strict dominance relationship between the $\mathcal{L}$ matrices of Regions 1 and 6. In general, however, if the patient is listed in Region 6, then she receives more frequent and higher quality liver offers. To quantify the magnitude of the violation of the dominance criterion 5.2, we define the following metric:

$\epsilon_5 = \max_h \left\{0, \sum_{j=k}^{L+1} [\mathcal{L}_1(j|h) - \mathcal{L}_6(j|h)]\right\}, 1 \leq k \leq L+1$, where $\mathcal{L}_1(j|h)$ and $\mathcal{L}_6(j|h)$ are the liver transition probability matrices of Region 1 and Region 6, respectively. ϵ_5 equals 0.0087 for the problems in Figure 6.5. As can be seen from Figure 6.5, the optimal control limits in Region 1 are higher than those of Region 6. This result is intuitive because as the patient receives more frequent liver offers with higher quality, she will be more selective.

6.4 CONCLUSIONS

This chapter considers the problem of accepting or declining a cadaveric liver offer. We derive structural properties of the MDP model, including conditions that guarantee the existence of a health-based and a liver-based control-limit policy. The conditions that ensure the

(a)Annual discount rate=0.99 (b)Annual discount rate=1

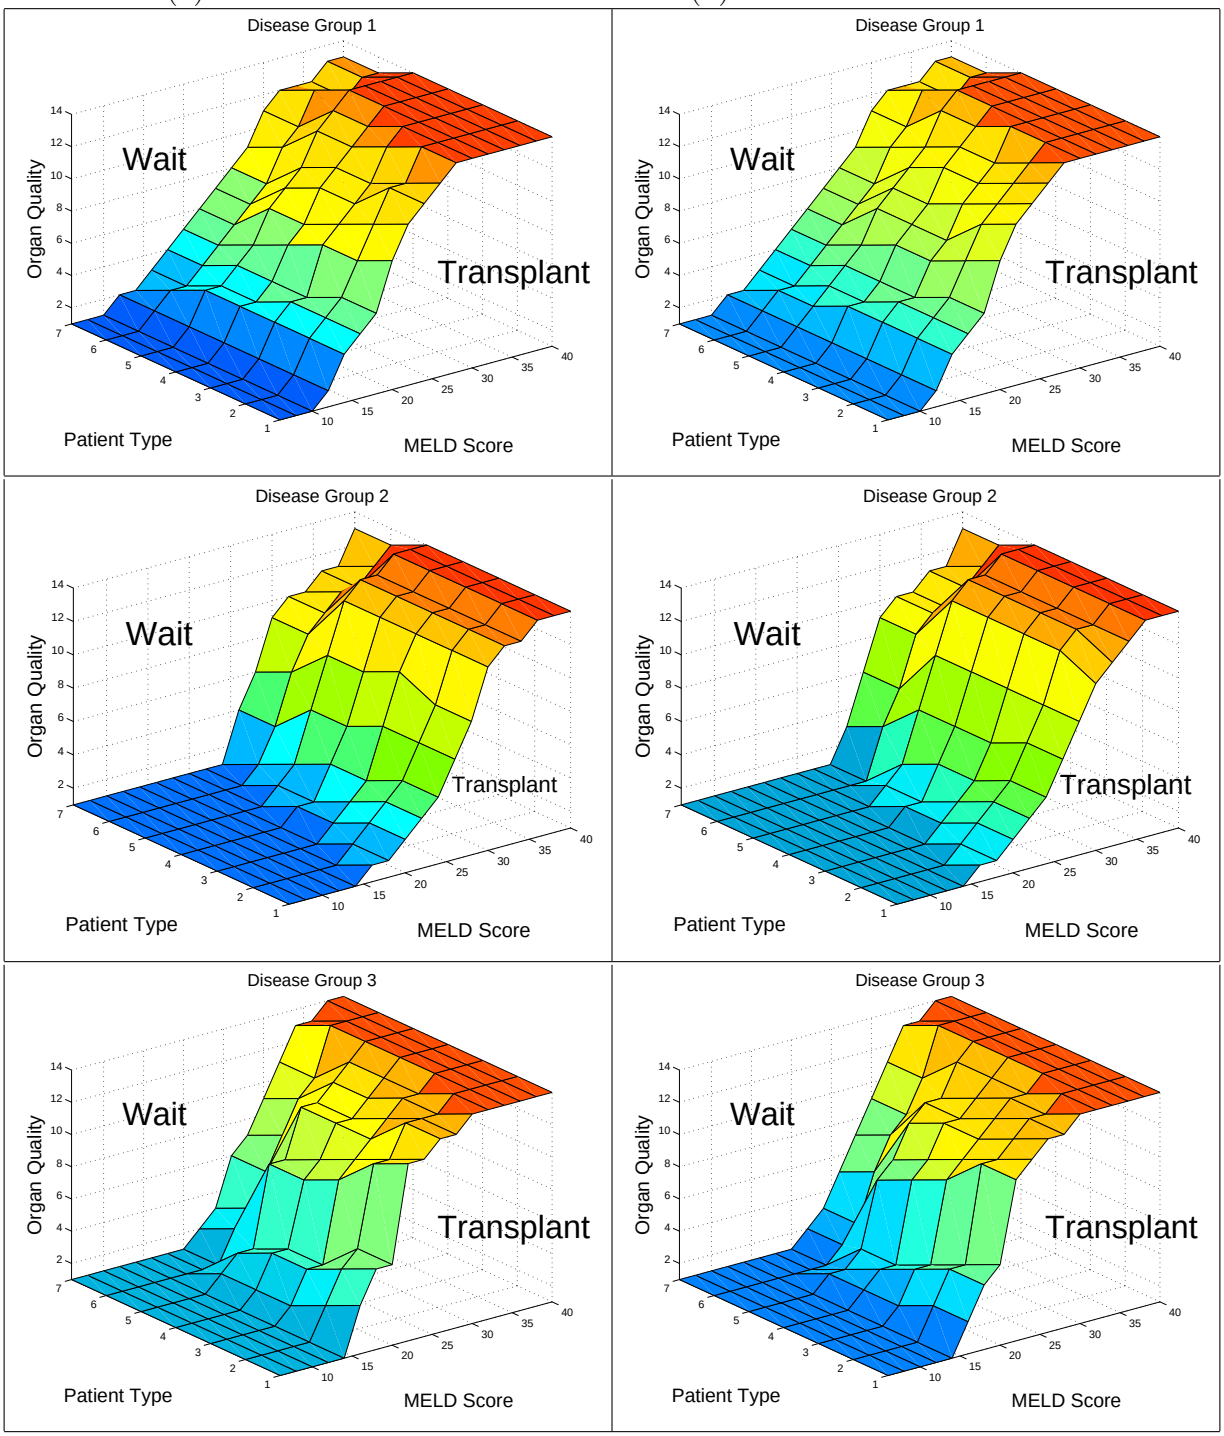

Figure 6.4: Transplant-Wait Decisions for Discounted and Undiscounted Problems

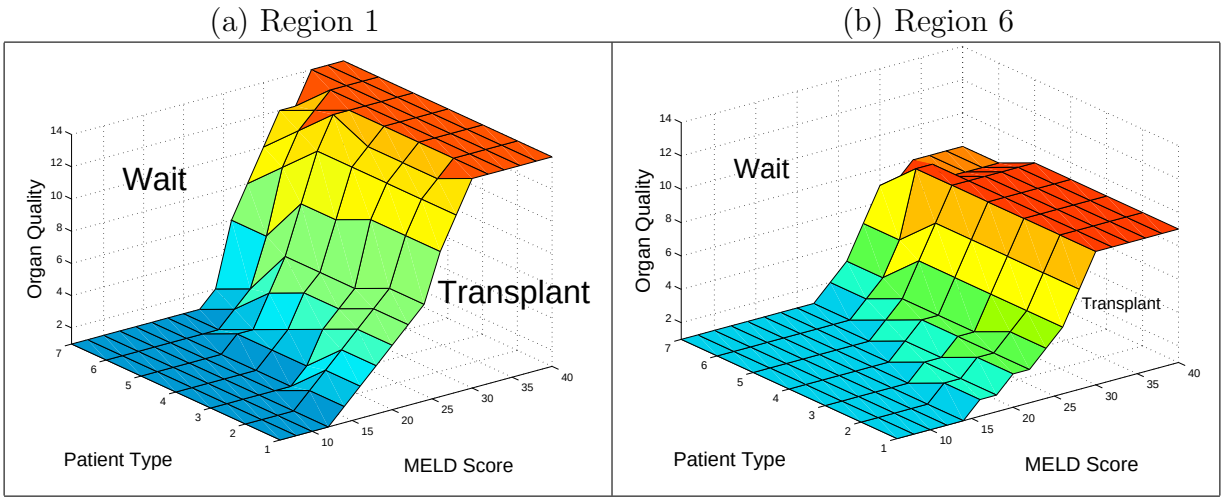

Figure 6.5: Transplant-Wait Decisions for Various Regions

existence of a health-based control-limit policy are stronger than those that guarantee the existence of a liver-based control-limit policy. We also compare the optimal control limits for the same patient listed in two different regions. We prove that if the patient is listed in a region where she receives more frequent and higher quality liver offers than another region, then the optimal control limits obtained when she is listed in the former region are higher than those obtained when she is listed in the latter region.

In all of our computational tests, the optimal policy is of liver-based control-limit type. All of the test problems satisfy the conditions of Theorem 6.3. However, some optimal policies are not of health-based control-limit type. In some examples, as the patient gets sicker, the probability of receiving a better liver increases significantly. In such cases, it is optimal to decline a liver offer in some health states even if it is optimal to accept that particular liver offer in better health states. Our computational tests also show that the location of the patient has a significant effect on liver offer probabilities and optimal control limits.

7.0 LIVING-AND-CADAVERIC-DONOR MODEL

In this chapter, we consider the general model presented in Chapter 3. This chapter is organized as follows. In Section 7.1, we derive structural properties of the MDP model and its optimal policy. In Section 7.2, we present and discuss computational results. We draw some conclusions and discuss future research directions in Section 7.3.

7.1 STRUCTURAL PROPERTIES

In this section, we derive some structural properties of the living-and-cadaveric-donor model (LCDM) given by 3.1. The following assumptions are common to all theorems:

(As1): The function $r(h, \ell, T_C)$ is nonincreasing in both h and ℓ, that is, as the patient gets sicker and/or the liver quality drops, her post-cadaveric-transplant reward does not increase.

(As2): The function $r(h, W)$ is nonincreasing in h, that is, as the patient gets sicker the intermediate reward does not increase.

(As3): The functions $r(h, \ell_{LD}, T_{LD})$ and $r'(h, \ell_{LD}, T_{LD})$ are nonnegative and nonincreasing in h, that is, as the patient gets sicker, her (net) post-living-donor-transplant reward does not increase.

The proof of the following theorem is similar to the proof of Theorem 6.1 and is omitted.

Theorem 7.1. *$V(h, \ell)$ is monotonically nonincreasing in ℓ, $\ell \in S_L$, $\forall h \in S_H$.*

Definition 7.1. *An* at-most-two-region liver-based policy *(AM2RL) for a particular health state is of the following form: For a given health state h, "Accept" the cadaveric organ if and only if the offered organ is of types $1, 2, \ldots, i(h)$, for some liver type $i(h)$, called the* AM2RL

control limit. *For all remaining health states, if* $a^*(h, i(h)+1) =$'W' *then* $a^*(h,\ell) =$'W' *holds for* $\ell = i(h)+2, \ldots, L$ *and if* $a^*(h, i(h)+1) =$'T_{LD}' *then* $a^*(h,\ell) =$'T_{LD}' *holds for* $\ell = i(h)+2, \ldots, L$.

An example for AM2RL policy is presented in Figure 7.1(c). Theorem 7.2 proves that *(As1)*, *(As2)* and *(As3)* suffice to ensure the existence of an optimal AM2RL policy. We first state two lemmas whose proofs are obvious and are omitted.

Lemma 7.1. *For a given health state* h, *if* $a^*(h,\ell)$='T_{LD}', *then* $a^*(h,\ell')$='T_{LD}' *for* $\ell' > \ell$.

Lemma 7.2. *For a given health state* h, *if* $a^*(h,\ell)$='W', *then* $a^*(h,\ell')$='W' *for* $\ell' > \ell$.

Theorem 7.2. *There exists an optimal AM2RL policy.*

Proof. We first show that $a^*(h,\ell+1)$='$\mathrm{T_C}$' implies $a^*(h,\ell)$='$\mathrm{T_C}$', $\forall h \in S_H$. If $a^*(h,\ell+1)$='$\mathrm{T_C}$' for a given $h \in S_H$, then it is obvious that $r(h,\ell,T_C) \geq r(h,\ell{+}1,T_C) \geq r'(h,\ell_{LD},T_{LD})$. The following are also true about the optimality equations:

$$V(h,\ell+1) = r(h,\ell+1,T_C) \geq r(h,W) + \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h',\ell'|h,\ell+1) V(h',\ell')$$

and

$$V(h,\ell) = \max \left\{ r(h,\ell,T_C), r(h,W) + \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h',\ell'|h,\ell) V(h',\ell') \right\}.$$

Because

$\lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h',\ell'|h,\ell+1) V(h',\ell') = \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h',\ell'|h,\ell) V(h',\ell')$ holds in the above expressions, as a result of *(As1)* we can write the following:

$$r(h,\ell,T_C) \geq r(h,\ell+1,T_C) \geq r(h,W) + \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h',\ell'|h,\ell) V(h',\ell'),$$

which is equivalent to $a^*(h,\ell)$='$\mathrm{T_C}$', from which the first part follows. Let $a^*(h,\ell)$='$\mathrm{T_C}$' for $\ell = 1, \ldots, \ell^*$ and $a^*(h,\ell^*+1) \neq$ '$\mathrm{T_C}$' . As a result of Lemma 7.1 and Lemma 7.2, the optimal policy will be either 'W' or '$\mathrm{T_{LD}}$' for $\ell > \ell^*$, from which the result follows. □

Theorem 7.3 proves the monotonicity of $V(h,\ell)$ in h under a set of intuitive conditions. The proof of this theorem is similar to Theorem 6.2 and is omitted.

Theorem 7.3. *Let $\mathcal{H}$ be an IFR matrix and (6.2) hold for all h and ℓ. Then $V(h,\ell)$ is also nonincreasing in h.*

Definition 7.2. *An* at-most-three-region (AM3R) policy *satisfies the following two conditions: For a given liver type ℓ, choose the "Wait" action if and only if the observed health state is one of the states $1, 2, \ldots, j(\ell)$, for some health state $j(\ell)$. Furthermore, there exists an AM2RL policy for each health state h.*

An example for AM3R policy is presented in Figure 7.1(c). The following theorem, the main result of this section, gives a set of intuitive conditions that ensure the existence of an optimal AM3R policy. Note that the conditions of this theorem and Theorem 5.3 have similar interpretations.

Theorem 7.4. *Let $\mathcal{H}$ be an IFR matrix and suppose $\mathcal{L}$ satisfies (6.2), and $r(h,\ell,T_C)$, $r'(h,\ell_{LD},T_{LD})$ and $\mathcal{H}$ satisfy (6.6) and (6.7). Then there exists an optimal AM3R policy.*

Proof. Note that Theorem 7.2 applies under the above assumptions. Therefore, if we show that for any ℓ, $a^*(h,\ell)$ ='W' implies $a^*(h-1,\ell)$ ='W', the result will then follow.

Note also that the monotonicity results in Theorems 7.1 and 7.3 hold. Now assume that for a given ℓ and h, $a^*(h,\ell)$ ='W' but $a^*(h-1,\ell)$ is uniquely either '$\mathrm{T_{LD}}$' or '$\mathrm{T_C}$'. Then, there are two cases:

Case 1. $a^*(h-1,\ell)$='$\mathrm{T_{LD}}$'

Because $\mathcal{H}$ is IFR, condition (5.8) is satisfied, condition (6.7) implies condition (5.9), the proof of Theorem 5.3 applies to this situation. As a result, $a^*(h-1,\ell)$ must also be 'W' from which the result follows.

Case 2. $a^*(h-1,\ell)$='$\mathrm{T_C}$'

Because $\mathcal{H}$ is IFR and conditions (6.6) and (6.7) are satisfied, the proof of Theorem 6.4 applies to this situation. As a result, $a^*(h-1,\ell)$ must also be 'W' from which the result follows. □

Theorem 7.5 compares the optimal policies of two identical patients who have different disutility functions. Namely, if a patient has a higher disutility of using the living donor than another patient, then the number of states in which "Transplant the living-donor liver"

action is optimal is smaller than or equal to the number of states that the second patient chooses to "Transplant the living-donor liver".

Theorem 7.5. *Let Π_1 and Π_2 be two instances that have identical $\mathcal{P}, r(h,W), r(h,\ell_C,T_C)$ and $r(h,\ell_{LD},T_{LD})$. Let $V_1(h,\ell)$ and $V_2(h,\ell)$ be the optimal value functions and $\rho_1(h)$ and $\rho_2(h)$ be the disutility functions of Π_1 and Π_2, respectively. Let $\rho_1(h) = \rho_1 r(h,\ell_{LD},T_{LD})$ and $\rho_2(h) = \rho_2 r(h,\ell_{LD},T_{LD})$, where $\rho_2 = \alpha\rho_1$ and $0 \leq \alpha \leq 1$. Then the following hold:*

(a) $V_1(h,\ell) \leq V_2(h,\ell)$, for $h \in S_H$ and $\ell in S_L$.

(b) Let $a_1^*(h,\ell)$ and $a_2^*(h,\ell)$ be the optimal actions of Π_1 and Π_2 for state (h,ℓ), respectively. If $a_1^*(h,\ell)$='T_{LD}', then $a_2^*(h,\ell)$='T_{LD}' must also hold.

Proof. (a) For any h and ℓ, we can write the following as a result of (3.1):

$$V_1(h,\ell) = \max\left\{r(h,\ell_{LD},T_{LD}) - \rho_1(h), r(h,\ell,T_C), r(h,W) + \lambda\sum_{h'}\sum_{\ell'}\mathcal{P}(h',\ell'|h,\ell)V_1(h',\ell')\right\}$$

and

$$V_2(h,\ell) = \max\left\{r(h,\ell_{LD},T_{LD}) - \rho_2(h), r(h,\ell,T_C), r(h,W) + \lambda\sum_{h'}\sum_{\ell'}\mathcal{P}(h',\ell'|h,\ell)V_2(h',\ell')\right\}.$$

Because $\rho_1(h) \geq \rho_2(h)$, $r(h,\ell_{LD},T_{LD}) - \rho_1(h) \leq r(h,\ell_{LD},T_{LD}) - \rho_2(h)$ holds. Since the second and third components of the above equations are identical, $V_1(h,\ell) \leq V_2(h,\ell)$ follows.

(b) If we show that $r_2'(h,\ell_{LD},T_{LD}) = r(h,\ell_{LD},T_{LD}) - \rho_2(h) = [r(h,\ell_{LD},T_{LD}) - \rho_1(h)](1+\alpha') = r_1'(h,\ell_{LD},T_{LD})(1+\alpha')$ for some $\alpha' > 0$, then the proof of Part (a) of Theorem 5.4 applies to this situation and the result follows.

$$\begin{aligned} r_2'(h,\ell_{LD},T_{LD}) &= r(h,\ell_{LD},T_{LD})[1-\rho_2] \\ &= r(h,\ell_{LD},T_{LD})[1-\alpha\rho_1] \\ &= \frac{r_1'(h,\ell_{LD},T_{LD})}{[1-\rho_1]}[1-\alpha\rho_1] \\ &= r_1'(h,\ell_{LD},T_{LD})[1+\alpha'], \end{aligned} \tag{7.1}$$

where (7.1) follows because $\frac{[1-\rho_1]}{[1-\alpha\rho_1]} \geq 1+\alpha'$ for some $\alpha' \geq 0$ by definition of α, from which the result follows. □

The LCDM generalizes the CDM by considering the availability of a living-donor liver. We show that without making any additional assumptions, an LCDM problem can be converted into an equivalent CDM problem. This conversion is accomplished by treating the living-donor liver as a cadaveric liver offer and changing the $r(h, \ell, T_C)$ function. This transformation is based on the fact that the patient would never accept a cadaveric liver that results in a lower total expected discounted post-transplant reward than the living-donor liver. Similarly, the patient would never accept the living-donor liver if doing so results in a lower total expected discounted post-transplant reward than transplanting a cadaveric-donor liver. Let $\tilde{r}(h, \ell, T_C)$ be the post-transplant reward function of the transformed problem. We obtain $\tilde{r}(h, \ell, T_C)$ as follows. For any h, if there exists an $\ell^*(h)$ such that $r(h, \ell^*(h), T_C) \leq r'(h, \ell_{LD}, T_{LD})$ and $r(h, \ell^*(h) - 1, T_C) > r'(h, \ell_{LD}, T_{LD})$, then $\tilde{r}(h, \ell(h), T_C) = r'(h, \ell_{LD}, T_{LD})$ for $\ell(h) = \ell^*(h), \ldots, L+1$ and $\tilde{r}(h, \ell(h), T_C) = r(h, \ell(h), T_C)$ for $\ell(h) < \ell^*(h), \ell(h) \in S_L$. Theorem 7.6 proves that the optimal value function obtained by solving the transformed problem is equal to the optimal value function of the original problem. Corollary 7.1 shows how the optimal actions of the original problem can be obtained from those of the transformed problem. The proof of the corollary is obvious and is omitted.

Theorem 7.6. *Let $\tilde{V}$ be the optimal value function of the transformed problem. Then $V(h, \ell) = \tilde{V}(h, \ell), h \in S_H, \ell \in S_L$.*

Proof. Note that we can rewrite the optimality equation for the LCDM as follows:

$$V(h, \ell(h)) = \max \left\{ r(h, \ell(h), T_C), r(h, W) + \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h', \ell' | h, \ell(h)) V(h', \ell') \right\}, \\ h \in S_H, \ell(h) = 1, \ldots, \ell^*(h) - 1, \quad (7.2)$$

$$V(h, \ell(h)) = \max \left\{ r'(h, \ell_{LD}, T_{LD}), r(h, W) + \lambda \sum_{h'} \sum_{\ell'} \mathcal{P}(h', \ell' | h, \ell(h)) V(h', \ell') \right\}, \\ h \in S_H, \ell(h) = \ell^*(h), \ldots, L+1, \quad (7.3)$$

where (7.2) follows because $r(h, \ell(h), T_C) \geq r(h, \ell_{LD}, T_{LD})$ for $\ell(h) < \ell^*(h)$ and (7.3) follows because $r(h, \ell(h), T_C) \leq r(h, \ell_{LD}, T_{LD})$ for $\ell(h) \geq \ell^*(h)$ by definition of $\ell^*(h)$. Combining

(7.2) and (7.3) results in the following set of recursive equations:

$$V(h,\ell) = \max\left\{\tilde{r}(h,\ell,T_C), r(h,W) + \lambda \sum_{h'}\sum_{\ell'} \mathcal{P}(h',\ell'|h,\ell)V(h',\ell')\right\}, \quad h \in S_H, \ell \in S_L,$$

which is identical to the optimality equations of the transformed problem, from which the result follows. □

Corollary 7.1. *Let $\tilde{a}^*(s)$ be the optimal decision for the transformed problem when the state is s. Then, for any h, if $\tilde{a}^*(h,\ell(h))$='T_C' for $\ell(h) < \ell^*(h)$, then $a^*(h,\ell(h))$='T_C' for $\ell(h) < \ell^*(h)$. Similarly, if $\tilde{a}^*(h,\ell(h))$='T_C' for $\ell(h) \geq \ell^*(h)$, then $a^*(h,\ell(h))$='T_{LD}' for $\ell(h) \geq \ell^*(h)$. Also if $\tilde{a}^*(h,\ell(h))$='W' then $a^*(h,\ell(h))$='W' for $\ell(h) \in S_L$.*

Note that the results in this section do not follow directly from the results in Section 6.2. However, we can use this transformation to convert an LCDM problem into a CDM problem and solve it as a CDM.

7.2 COMPUTATIONAL RESULTS

We solve the LCDM using real data. The estimation of parameters is described in Chapter 4 and Section 5.3. We use the same values for the health state space, health and liver transition probability matrices and the reward function as in the CDM described in Chapter 6. All of our computational experiments in this section use the national liver transition probability matrix. We are unaware of any studies that quantify the disutility of living-donor liver transplants. We define $\rho(h)$ as a linear function of $r(h, l_{LD}, T_{LD})$, i.e. $\rho(h)$= $\rho_0 \cdot r(h, l_{LD}, T_{LD})$. Then $r'(h, l_{LD}, T_{LD}) = (1-\rho_0) \cdot r(h, l_{LD}, T_{LD})$. This definition implies that the sicker the patient, the smaller the disutility associated with accepting the living-donor liver.

We consider the same patient considered in Chapters 5 and 6, a 60-year old female patient with primary biliary cirrhosis who has blood type A. We use a 0.99 annual discount rate (daily $\lambda = 0.999972$) and ρ_0=0.1. Figure 7.1(c) shows the optimal policy when the only living-donor liver available to this patient is Organ 2, whose donor characteristics are given in Table 5.2. As can be seen from the figure, as the patient gets sicker, the "Transplant

cadaveric liver" action is optimal for more cadaveric liver types and the "Wait" action is optimal only in healthier states. This policy is an example for the AM3R policy described in Section 7.1. We tested a total of 240 instances and all optimal policies were of AM3R type.

Figure 7.1 compares the optimal policy for LCDM with the optimal policies for the LDM and CDM for the same patient and living-donor liver when annual discount rate is 0.99 (λ=0.999972) and $\rho_0 = 0.1$. Note that the optimal policy obtained by LDM is a vertical line in this figure, because when the patient is not listed on the waiting list she is indifferent between various cadaveric liver offers. As expected, the "Transplant the living-donor liver" action occurs in sicker health states for the LCDM than for the LDM. Similarly, the patient becomes more selective in accepting the cadaveric liver offers for the LCDM.

As expected, the disutility function affects the optimal policies significantly. Figure 7.2 shows the optimal policies for various values of ρ_0. As the disutility associated with transplanting the living-donor liver increases, the number of states where "Transplant living-donor organ" action is optimal decreases.

As in the LDM, different living-donor livers may result in different optimal policies. Figure 7.3 shows the optimal policies for the same patient and disutility function and various living-donor livers when $\rho_0 = 0.1$ and annual discount rate equals 0.99. Organ types that are used in Figure 5.2 are defined in Table 5.2. As can be seen from the figure, as the quality of the organ drops the number of states in which it is optimal to "Transplant living-donor organ" decreases. This result is intuitive because as the quality of the living donor drops, the patient benefits less from the living-donor transplantation. As a result, the patient is less selective when considering the cadaveric liver offers.

As in the LDM and CDM, different diseases and different discount factors have different control limits. Figure 7.4 shows the optimal policies for different diseases and discount factors when the only available living-donor liver to the patient is Organ 2 and $\rho_0 = 0.1$. Disease groups in Figure 7.4 are defined in Table A.1. As demonstrated in the figure, the optimal policies for Disease Group 3 have the largest "Wait" regions. Our results also show that the undiscounted problem have a larger "Wait" region than the discounted problem.

As can be seen from Figure 7.4, when the patient is in disease group 2, the optimal policy suggests the following: she should accept a cadaveric liver offer that is of liver type

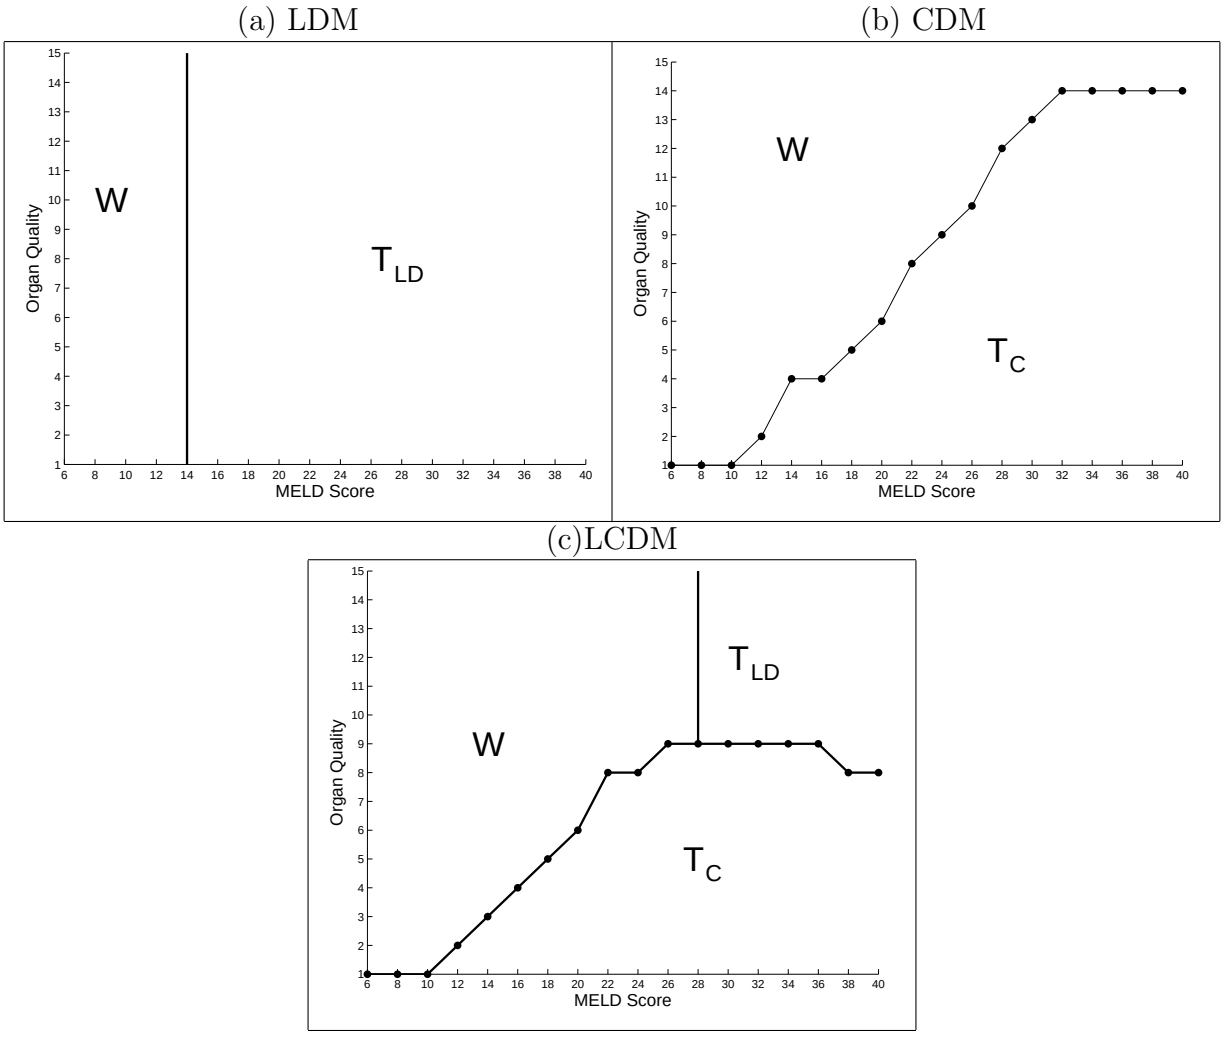

Figure 7.1: Transplant-Wait Decisions for Different Models

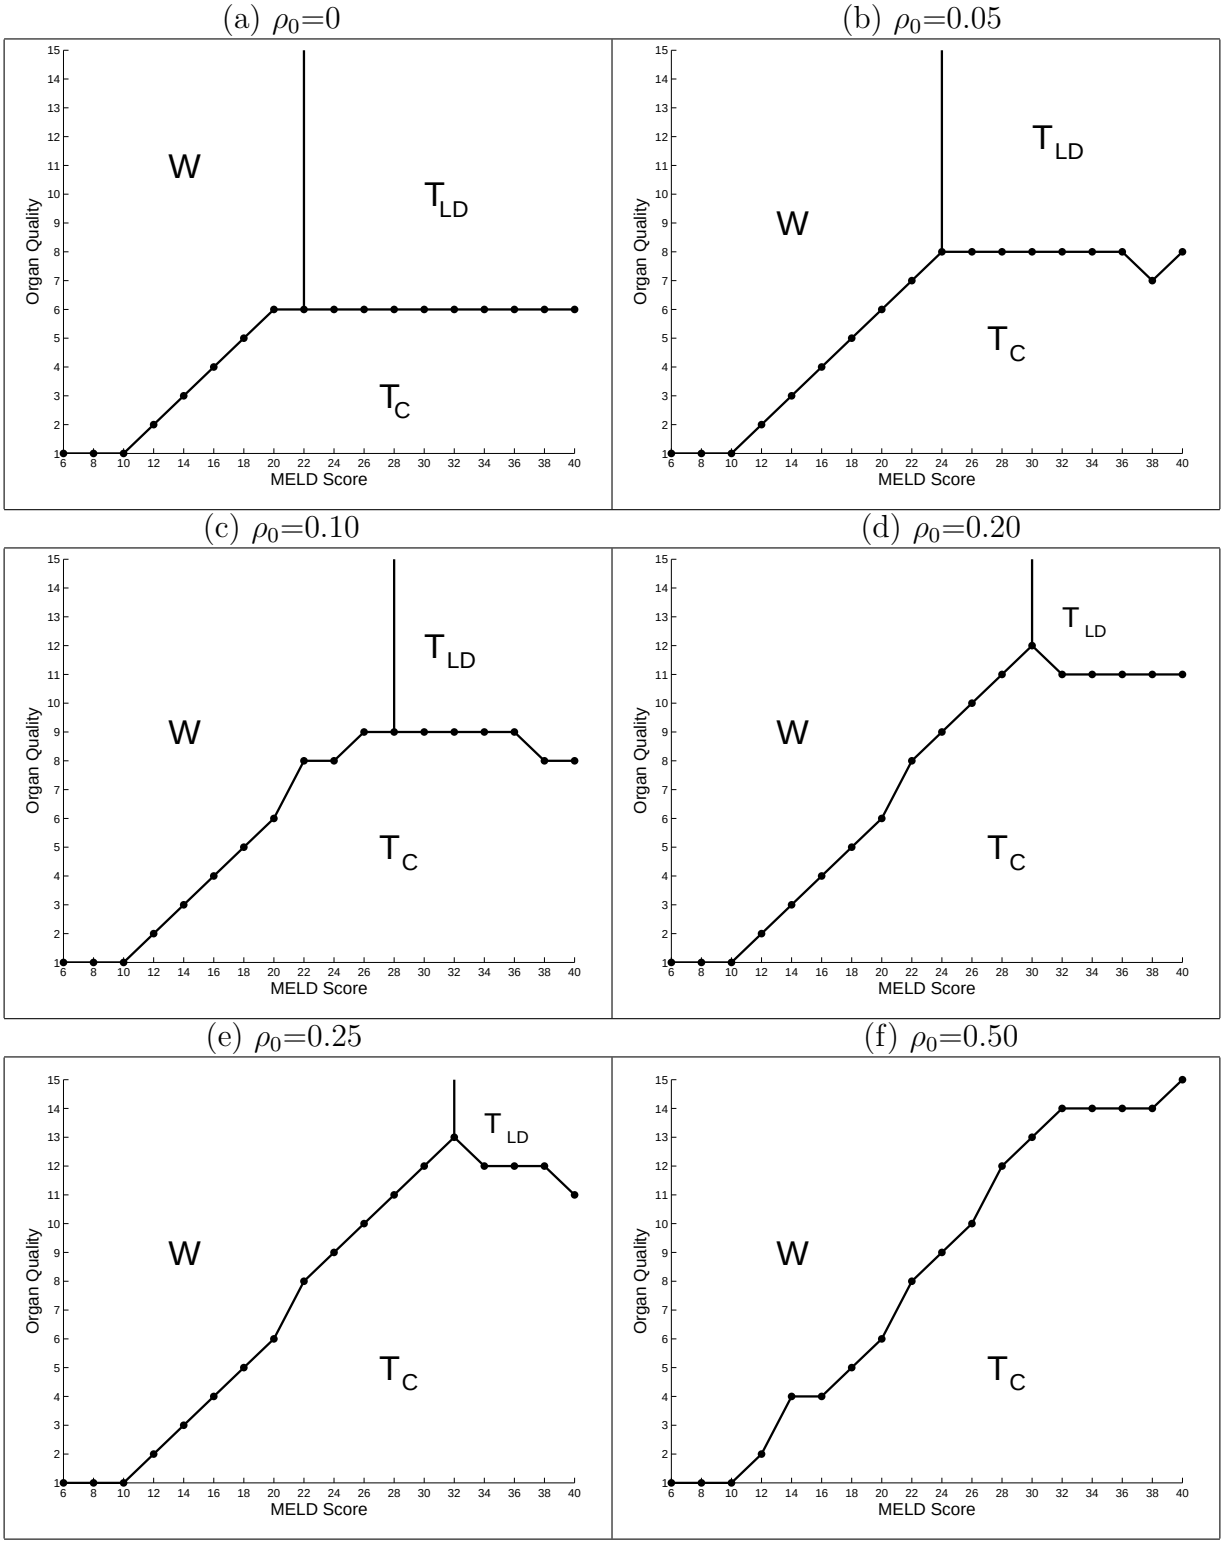

Figure 7.2: Transplant-Wait Decisions for Different Disutility Functions

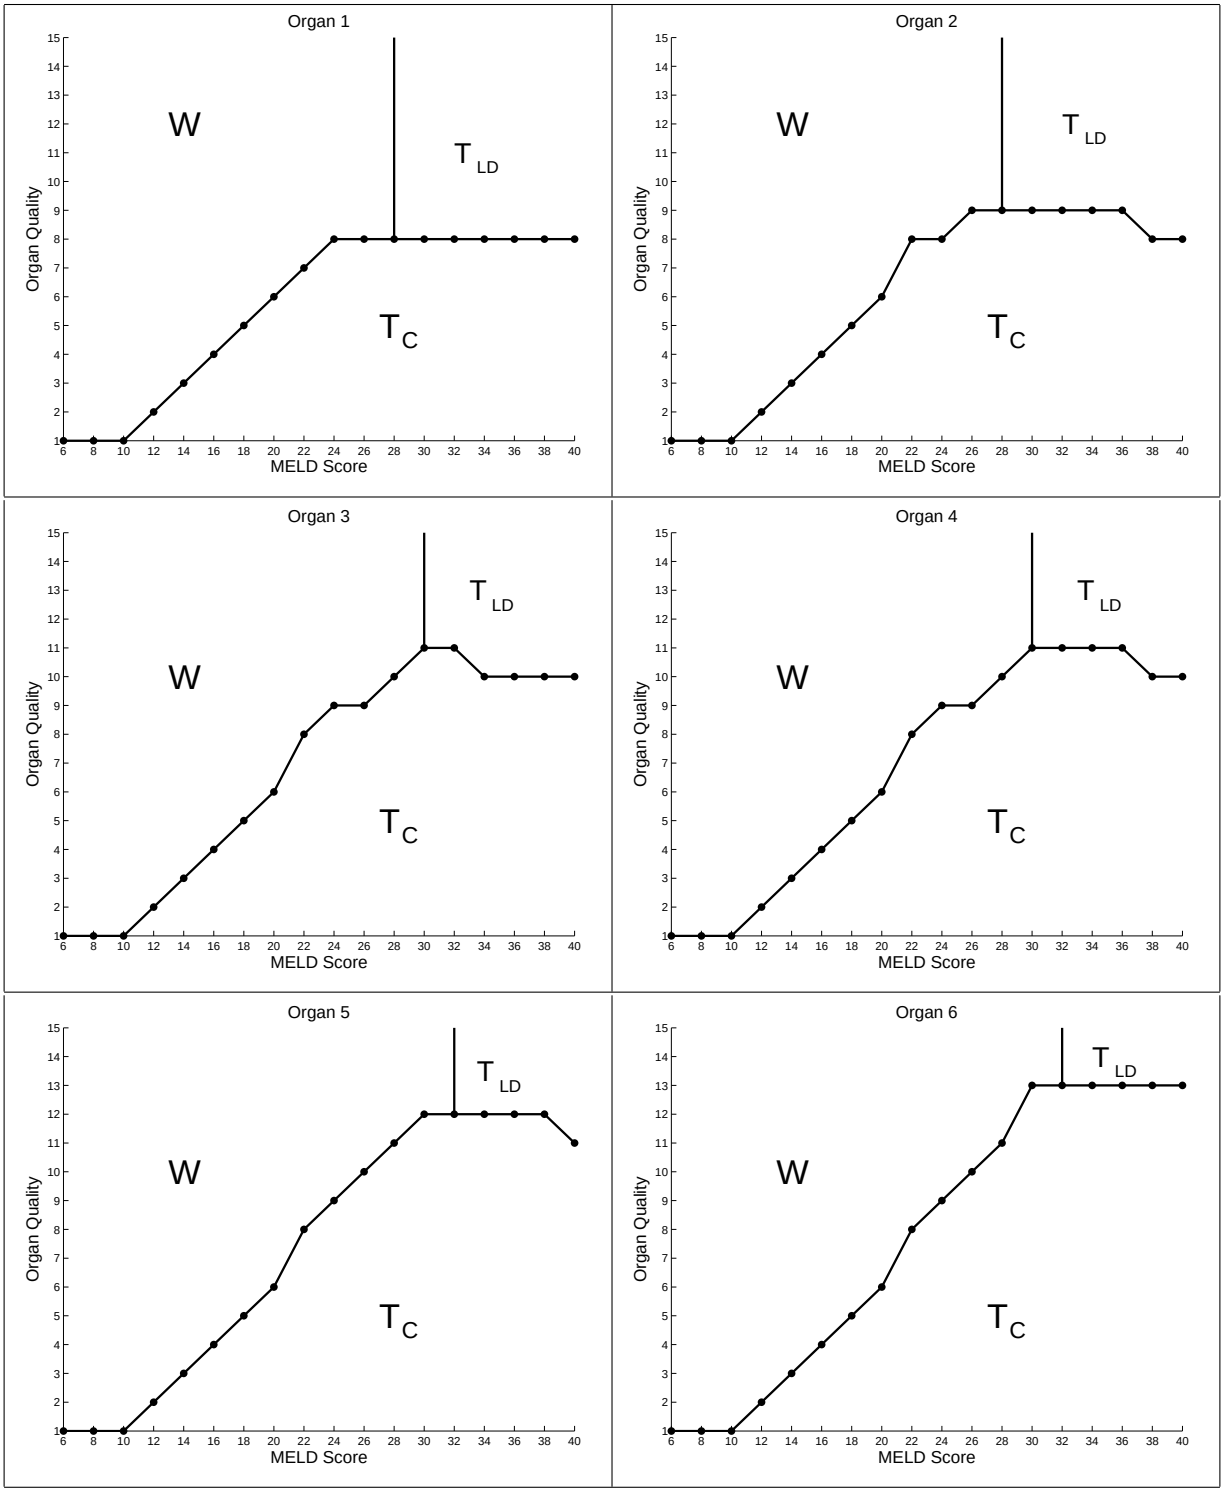

Figure 7.3: Transplant-Wait Decisions for Various Organs

9 when her MELD score is 34 and 38, and she should decline the same liver type when her MELD score is 36. This result may appear counterintuitive, however, recall that the current definition of disutility function does not ensure the existence of a complete ordering between living-donor and cadaveric liver types. For this particular example, the transplantation of the living-donor liver results in a smaller net total expected discounted reward than the cadaveric liver type 9 transplant when the patient's MELD score is 34 or 38. On the other hand, the living-donor liver transplantation results in a larger net total expected discounted reward than the cadaveric liver type 9 transplant when the patient's MELD score is 36. This is an example for the nonexistence of a complete ordering of the liver types. It is also possible that our estimation for the total post-transplant reward is not accurate, which indicates that the MELD system may not be a good measure of the total post-transplant life expectancy.

7.3 CONCLUSIONS

This chapter considers the most realistic decision problem faced by patients who have a living-donor liver available: should an offered cadaveric organ of a given quality be accepted or declined? If declined, should the living-donor liver be transplanted? The model presented in this chapter extends the one in Chapter 5 by combining the two models that are described in Chapters 5 and 6.

We derive structural properties of the MDP model, including conditions that guarantee the existence of an AM3R policy. In all of our computational tests, the optimal policy is of AM3R type. Our computational tests also show that the disutility associated with transplanting the living-donor liver significantly affects the optimal policy. In general, as the disutility function increases, the number of states in which it is optimal to "Wait" and to "Transplant the cadaveric liver" increases and the number of states where it is optimal to "Transplant the living-donor liver" decreases.

(a)Annual discount rate=0.99 (b)Annual discount rate=1

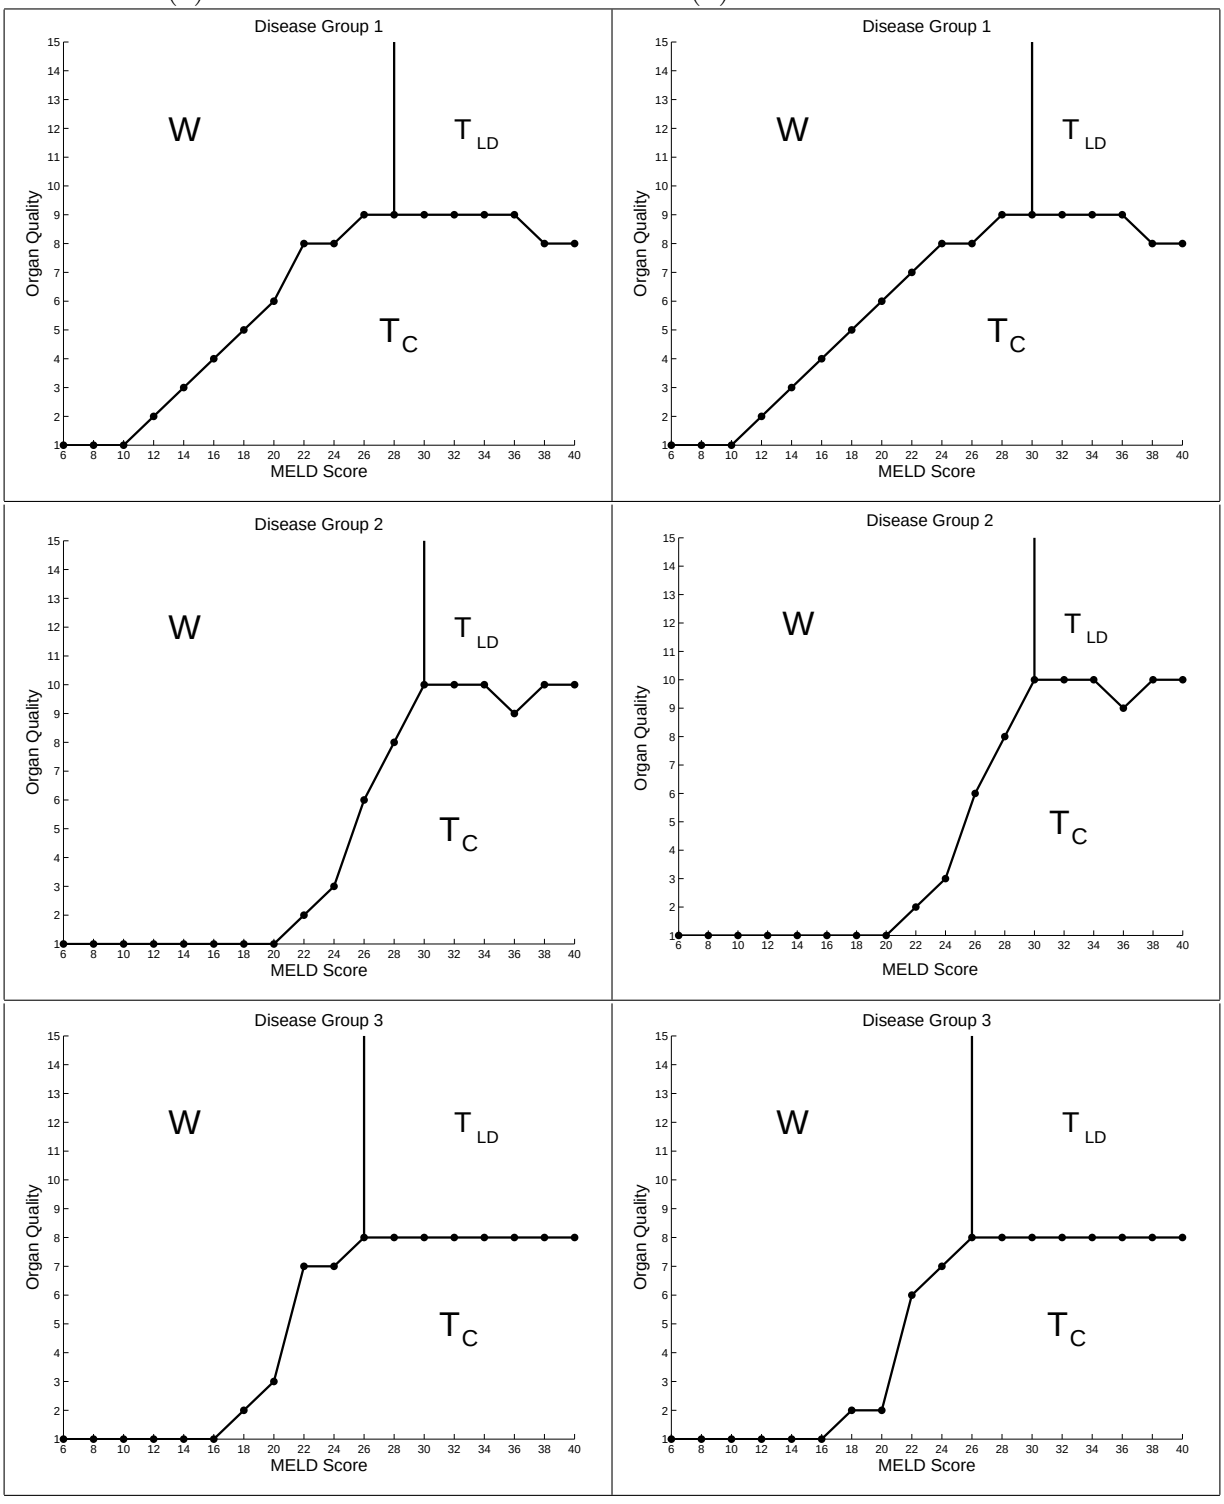

Figure 7.4: Transplant-Wait Decisions for Discounted and Undiscounted Problems

8.0 SUMMARY AND FUTURE RESEARCH

8.1 SUMMARY

This dissertation focuses on the decision problem faced by patients with ESLD: which liver offer to accept and which to refuse? Previous research [40, 41, 142, 165, 166] has focused on the optimal allocation of the organs to the patients. However, such studies must consider many political and ethical issues such as equity among various groups. Because this dissertation views and solves the problem from the patient's perspective, it does not need to consider such issues and therefore the resulting optimal policies are immediately implementable. Furthermore, unlike previous researchers, we employ detailed models of patient health and organ quality, and use real data to solve the decision problems. As a result, the models in this dissertation are not limited by unrealistic assumptions.

The most general model considers the problem of determining whether a patient should accept a cadaveric liver offer, or decline the cadaveric liver offer and accept the living-donor liver, or decline both cadaveric and living-donor livers. We formulate this problem as an infinite horizon discrete-time MDP model. Chapters 5 and 6 consider two specific cases of this MDP model: the living-donor-only model, in which we assume that the patient has a living donor and is not listed on the waiting list; the cadaveric-donor-only model, in which the patient is assumed to be listed on the liver waiting list and does not have any living-donor liver(s). To the best of our knowledge, the living-donor-only model is the first model that considers the optimal timing of living-donor organ transplants. In all three models, we derive the structural properties of the MDP model, including sets of sufficient conditions that ensure the existence of control-limit policies.

These decision models could be generalized to address the problem of optimally timing other organ transplants, given the following components: a proxy for the patient health, the transition probabilities between patient health and organ offer states, a measure for organ quality, and a reward function for a given organ and patient health characteristics. If there are additional treatment options (such as dialysis in the kidney transplantation), then the action space and the state space could be expanded to include new actions.

8.2 FUTURE RESEARCH

There are several possible extensions to this dissertation. Section 8.2.1 discusses various approaches to modeling the effect of the waiting list. We present the problem of optimally choosing the regions for listing in Section 8.2.2. Finally, Section 8.2.3 describes several other future research directions.

8.2.1 Modeling the Waiting List

Under the current UNOS policy, the composition of the waiting list and the actual position of the patient on this list have a significant impact on the liver offer(s) that the patient receives. Therefore, the waiting list must be considered in the cadaveric-donor liver transplant decision problem. Neither the LCDM nor the CDM includes the waiting list in the decision model as part of the state space nor do they make any attempt to infer the patient's position on the waiting list. Instead, these models incorporate the effect of the composition of the waiting list and the patient's position on the list into the state transition probabilities. Considering the waiting list requires a more complicated model.

One possible method to model the waiting list is to assume that a full description of the waiting list is available to the patient at all times and to include the description of the waiting list as part of the state space. An optimal solution to this problem gives the best policy that the patient can attain.

However, under the current UNOS policy, information about the position of the patient and the composition of the waiting list are not fully available. Instead, the patient can learn

only partial information about the composition of the waiting list such as the current medical status or MELD scores of other patients in her OPO and region within some ranges [148]. Furthermore, throughout her waiting time, the patient receives information that is inherently related to the overall composition of the waiting list, but does not reveal the composition of the list with certainty. Based on this information, i.e. the stream of livers offered or not offered, the succession of health levels and partial information about the waiting list, the patient can make an inference as to the composition of the list and make accepting/declining decisions based on this inference. For instance, consider a patient in the Pittsburgh OPO who has a blood type of O and has a MELD score of 18 as of January 13, 2004. This patient can look at the UNOS web site for her OPO region and learn about the total number of people waiting in her OPO and partial information about their health status. As of January 13, 2004, this patient learns from the web site that there are 6 patients with MELD scores larger than 18 or Status 1 and there are 19 people having MELD scores between 11 and 18 [148]. Based on this information, this patient can make an estimation of her rank by assuming a probability distribution for the other patients' individual MELD scores. As of January 13, 2004, this particular patient is most probably at the top of the patients having MELD scores between 11 and 18.

Viewing this decision problem in this manner naturally leads to a partially observed Markov decision process (POMDP) formulation. It can be shown [112] that the POMDP can be formulated as a Markov decision process with an enlarged state space, namely the space of probability distributions over the underlying states. In the context of the problem considered here, the underlying or "core" states may be taken to be the composition of the waiting list. Considering the POMDP problem enables us to measure the price of hiding information from the patients.

8.2.2 Multiple Listing Problem

According to the current UNOS policy, a patient can be listed in more than one OPO at the same time. Although multiple listing may increase the probability of receiving more frequent and higher quality liver offers, patients do not list in all OPOs due to financial and

geographical constraints. Furthermore, there may not be a complete ordering of the OPOs. For instance, there is not a strict relationship between an OPO with more frequent but lower quality livers and another OPO with less frequent but higher quality livers. Similarly, the frequency and the quality of liver offers depend on the patient health. Listing in one OPO may result in more frequent and higher quality liver offers than another OPO for some health states. However, the second OPO may be preferable to the first one for some other health states. An interesting problem is to determine the optimal set of OPOs that will maximize the total expected reward for the patient. This problem can be formulated and solved by combining a combinatorial optimization model with an MDP model.

8.2.3 Other Future Research Directions

In all of our computational tests, we use the total discounted expected-life days (ELD) as our reward function. However, it is well known that total discounted quality-adjusted life days (QALD) of the patient is a better measure for outcome in medical research [68]. Using QALD as the performance measure will not change the MDP model, however, it may have a significant effect on the optimal policies. A possible future research could be to consider QALD as the reward function and compare the resulting optimal policies with those that are obtained when ELD is used.

All of our models assume that the decision maker is both indifferent to the timing of the resolution of uncertainty and risk-neutral. However, these assumptions do not necessarily hold [32]. A possible future research would be to consider these and other patient preference issues, which may require a different reward function definition.

In this dissertation, we do not compare the optimal decisions obtained by using our models to the actual decisions made by the patients. Such a comparison enables us to measure the changes in the total (quality-adjusted) life expectancy of the patients when they use our decision models. It is also possible to measure the changes in the overall statistics of the organ allocation system such as liver refusal rates and the number of organs wasted when all patients use the decision models described in this dissertation, which requires the construction and execution of a simulation model.

APPENDIX A

CLASSIFICATION OF LIVER DISEASES

Table A.1: Classification of Liver Diseases [3, 127, 136]

Five-Group Classification	Ten-Group Classification	Category
1 Cirrhotic diseases	1	Primary biliary cirrhosis: primary biliary cirrhosis.
	2	Primary sclerosing cholangitis: Crohn's disease; ulcerative colitis; cholangitis with no bowel disease.
	3	Alcoholic liver disease: Lannec's cirrhosis.
	4	Autoimmune disorders: cirrhosis (drug or industrial exposure; cryptogenic; idiopathic); chronic autoimmune hepatitis (etiology unknown; postnecrotic).
2 Hepatitis infections	5	Hepatitis C and similar infections: postnecrotic cirrhosis (non-A, non-B; type C; type D; types B and C; types B and D); Lannec's cirrhosis (postnecrotic; other).
	6	Hepatitis B: postnecrotic cirrhosis (HBsAg-positive).
3 Acute liver diseases	7	Acute hepatic failure: acute hepatic necrosis (drug exposure; hepatitis A; hepatitis B, HBsAg-positive; non-A, non-B hepatitis; hepatitis C; hepatitis D; hepatitis B and C; hepatitis B and D; other acute viral infection); etiology unknown.
4 Cancers	8	Cancers: primary liver malignancy (hepatoma; fibrolamellar hepatocellular carcinoma; cholangiocarcinoma; hepatoblastoma; hemangioendothelioma; hemangiosarcoma; angiosarcoma).
5 Other liver diseases	9	Metabolic disorders: alpha-1-antitrypsin deficiency; glycogen storage disease type I; glycogen storage disease type II; hemochromatosis; hemosiderosis; hyperlipidemia type II; homozygous hypercholesterolemia; primary oxalosis or oxaluria; hyperoxaluria; tyrosinemia; Wilson's disease or other copper disorder; urea cycle disorder; Crigler-Najjar syndrome; Wolman's disease; protoporphyria; Niemann-Pick disease; abetalipoproteinemia; Gaucher's disease; Osler-Rendu-Weber disease; carbamoylphosphate synthase deficiency; amyloidosis; Wiskott-Aldrich syndrome.
	10	Other liver diseases: cirrhosis (postnecrotic hepatitis A); secondary biliary cirrhosis (Caroli's disease; choledochal cyst; other); familial cholestasis (Byler disease; other); cholestatic liver disease not listed above; neonatal hepatitis; biliary atresia (extrahepatic biliary atresia; hypoplasia; Alagille syndrome; other); congenital hepatic fibrosis; cystic fibrosis; Budd-Chiari syndrome; benign tumor (hepatic adenoma; polycystic liver disease; other); liver disease induced by total parenteral nutrition or hyperalimentation; graft-versus-host disease; trauma; biliary stricture or stenosis; idiopathic adult ductopenia; unknown.

APPENDIX B

NATIONAL CLINICAL OVERSIGHT COMMITTEE

Affiliations are listed at the time of membership.

John Fung, MD, PhD
Associate Professor of Surgery and Transplantation Research Director
University of Pittsburgh School of Medicine
Pittsburgh, PA 15213

Roger Jenkins, MD
Chief of Transplantation Beth Israel/Deaconess Medical Center
Harvard Medical School
Boston, MA 02215

John R. Lake, MD
Professor of Medicine and Surgery
University of Minnesota
Minneapolis, MN 55455

Wallis Marsh, MD
Associate Professor of Surgery
University of Pittsburgh School of Medicine
Pittsburgh, PA 15213

Andre C. Stieber, MD, FACS
Associate Professor of Surgery, Liver Transplantation
Emory University School of Medicine
Atlanta, GA 30322

Russell H. Wiesner, MD
Medical Director of Liver Transplantation
Mayo Clinic
Rochester, MN 55905

Rowen K. Zetterman, MD, FACP, FACG
Professor, Department of Internal Medicine
University of Nebraska Medical Center
Omaha, NE 68198

BIBLIOGRAPHY

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

[2] O. Alagoz, C. L. Bryce, A. J. Schaefer, D. C. Angus, and M. S. Roberts. Predicting the future health states of liver disease patients using empirical stochastic models. *Medical Decision Making*, 22(6):541, 2002.

[3] O. Alagoz, C. L. Bryce, A. J. Schaefer, C. H. Chang, D. C. Angus, and M. S. Roberts. Incorporating biological natural history in simulation models: Empiric estimates of the progression of end-stage liver disease, 2004. Submitted to *Medical Decision Making*.

[4] O. Alagoz, L. M. Maillart, A. J. Schaefer, and M. S. Roberts. Determining the optimal timing of living donor liver transplantation using a Markov Decision Process (MDP) model. *Medical Decision Making*, 22(6):558, 2002.

[5] O. Alagoz, L. M. Maillart, A. J. Schaefer, and M. S. Roberts. The optimal timing of living-donor liver transplantation, 2004. Forthcoming in *Management Science*.

[6] M. Aoki. Optimal control of partially observable Markovian systems. *Journal of The Franklin Institute*, 280(5):367–386, 1965.

[7] K. J. Astrom. Optimal control of Markov processes with incomplete state information. *Journal of Mathematical Analysis and Applications*, 10:174–205, 1965.

[8] M. A. Badri and J. A. Hillingsworth. Emergency room scheduling: A simulation model. In P. Luker, editor, *Proceedings of the 1992 Summer Computer Simulation Conference. Twenty Fourth Annual Computer Simulation Conference, SCS, San Diego, CA, USA*, pages 194–198, 1992.

[9] F. F. Baesler and J. A. Sepulveda. Multi-objective simulation optimization for a cancer treatment center. In B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, editors, *Proceedings of the 2001 Winter Simulation Conference, Piscataway, NJ*, pages 1405–1411, 2001.

[10] Ž. Bajzer, M. Marušić, and S. Vuk-Pavlović. Conceptual frameworks for mathematical modeling of tumor growth dynamics. *Mathematical and Computer Modelling*, 23(6):31–46, 1996.

[11] D. Banjevic, A. K. S. Jardine, V. Makis, and M. Ennis. A control-limit policy and software for condition-based maintenance optimization. *INFOR*, 39(1):32–50, 2001.

[12] R. E. Barlow and F. Proschan. *Mathematical Theory of Reliability*. John Wiley and Sons, New York, NY, 1965.

[13] R. E. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957.

[14] D. P. Bertsekas. *Dynamic Programming and Stochastic Control (Volumes 1 and 2)*. Athena Scientific, Belmont, MA, 2001.

[15] D. P. Bertsekas and J. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA, 1996.

[16] I. Bilbao, J. Figueras, L. Grande, M. Clèries, E. Jaurrieta, J. Visa, and C. Margarit. Risk factors for death following liver retransplantation. *Transplantation Proceedings*, 35(5):1871–1873, 2003.

[17] D. Blackwell. Discounted dynamic programming. *Annals of Mathematical Statistics*, 36(1):226–235, 1965.

[18] C. C. Branas, E. J. MacKenzie, and C. S. ReVelle. A trauma resource allocation model for ambulances and hospitals. *Health Services Research*, 35(2):489–507, 2000.

[19] L. Brotcorne, G. Laporte, and F. Semet. Ambulance location and relocation models. *European Journal of Operational Research*, 147(3):451–463, 2003.

[20] P. Bruns. Optimality of randomized strategies in a Markovian replacement model. *Mathematical Methods of Operations Research*, 56(3):481–499, 2003.

[21] C. L. Bryce, O. Alagoz, A. J. Schaefer, D. C. Angus, and M. S. Roberts. Modeling the progression of liver disease over time: A comparison of level and direction-based algorithms. *Medical Decision Making*, 23(6):506, 2003.

[22] J. A. Cairns. Health, wealth and time preference. *Project Apraisal*, March:31–40, 1992.

[23] B. Carton. Doctors, U.S. Government move closer to backing payment when organs are donated. *Wall Street Journal*, 2002. February 14, 2002.

[24] A. R. Cassandra. *Exact and Approximate Algorithms for Partially Observable Markov Decision Processes*. PhD dissertation, Brown University, 1998.

[25] A. Cayley. Mathematical questions with their solutions, No. 4528. *Educational Times*, 23(18), 1875.

[26] London Health Sciences Center, 2004. Available from http://www.lhsc.on.ca/ transplant/ethics.htm, information and data accessed on May 31, 2004.

[27] H. P. Chao. Inventory policy in the presence of market disruptions. *Operations Research*, 35(2):274–281, 1987.

[28] G.B. Chapman. Time discounting of health outcomes. In G. Loewenstein, D. Read, and R. F. Baumesiter, editors, *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*, New York, NY, 2003. Russell Sage Foundation.

[29] M. Chen and R. M. Feldman. Optimal replacement policies with minimal repair and age-dependent costs. *European Journal of Operational Research*, 98(1):75–84, 1997.

[30] F. Cheng and S. P. Sethi. A periodic review inventory model with demand influenced by promotion decisions. *Management Science*, 45(11):1510–1523, 1999.

[31] H. Cheng. *Algorithms for Partially Observable Markov Decision Processes*. PhD dissertation, University of British Columbia, 1988.

[32] S.H. Chew and J. L. Ho. Hope: An empirical study of attitudes toward the timing of uncertainty resolution. *Journal of Risk and Uncertainty*, 151:949–953, 1994.

[33] C. Chow and J. Tsitsiklis. The complexity of dynamic programming. *Journal of Complexity*, 5(4):466–488, 1989.

[34] R. Church, P. Sorensen, and W. Corrigan. Manpower deployment in emergency services. *Fire Technology*, 55(2):420–427, 2003.

[35] CNN, 2004. Available from http://www.cnn.com/2003/HEALTH/03/14/life.expectancy.ap/, information and data accessed on June 1, 2004.

[36] COD, 2004. Available from http://www.donatelife.net, information and data accessed on February 15, 2004.

[37] CONSAD Research Corporation. *An Analysis of Alternative National Policies for Allocating Donor Livers for Transplantation*. CONSAD Research Corporation, Pittsburgh, PA, 1995.

[38] D. R. Cox. Regression models in life tables (with discussion). *Journal of the Royal Statistical Society B*, 34:187–220, 1972.

[39] D. Coyle and K. Tolley. Discounting of health benefits in the parmacoeconomics analysis of drug therapies: An issue for debate? *PharmacoEconomics*, 2:153–162, 1992.

[40] I. David and U. Yechiali. A time-dependent stopping problem with application to live organ transplants. *Operations Research*, 33(3):491–504, 1985.

[41] I. David and U. Yechiali. One-attribute sequential assignment match processes in discrete time. *Operations Research*, 43(5):879–884, 1995.

[42] Y. De Serres. Simultaneous optimization of flow-control and scheduling in a single server queue with two job classes: Numerical results and approximation. *Computers and Operations Research*, 18(4):361–378, 1991.

[43] F. D'Epenoux. A probabilistic production and inventory problem. *Management Science*, 10(1):98–108, 1963.

[44] C. Derman. On sequential decisions and Markov chains. *Management Science*, 9(1):16–24, 1962.

[45] C. Derman. On optimal replacement rules when changes of state are Markovian. In R. Bellman, editor, *Mathematical Optimization Techniques*, pages 201–210, Berkeley, CA, 1963. The Rand Corporation.

[46] C. Derman. Optimal replacement and maintenace under Markovian deterioration with probability bounds on failure. *Management Science*, 9(3):478–481, 1963.

[47] J. L. Dienstag and K. J. Isselbacher. Chapter 295: Acute viral hepatitis. In E. Braunwald, A. S. Fauci, D. L. Kasper, S. L. Hauser, D. L. Longo, and J. L. Jameson, editors, *Harrison's Principles of Internal Medicine*, pages 1721–1736, New York, NY, 2001. McGraw-Hill.

[48] X. Ding, M. L. Puterman, and A. Bisi. The censored newsvendor and the optimal acquisition of information. *Operations Research*, 50(3):517–527, 2002.

[49] J. B. R. do Val and J. L. F. Salles. Optimal production with preemption to meet stochastic demand. *Automatica*, 35(11):1819–1828, 1999.

[50] K. Dudek, P. Nyckowski, K. Zieniewicz, B. Michalowicz, J. Pawlak, P. Malkowski, and M. Krawczyk. Liver retransplantation: Indications and results. *Transplantation Proceedings*, 34(2):638–639, 2002.

[51] E. B. Dynkin. Controlled random sequences. *Theory and Probability Applications*, 10:1–14, 1965.

[52] V. S. Elliott. Transplant scoring system called more fair. *American Medical News*, May 13, 2002. Available from http://www.ama-assn.org/amednews/2002/05/13/hll20513.htm.

[53] J. E. Falk, S. W. Palocsay, W. J. Sacco, W. S. Copes, and H. R. Champion. Bounds on a trauma outcome function via optimization. *Operations Research*, 40(1):586–595, 1992.

[54] E. A. Feinberg and A. B. Piunovskiy. Multiple objective nonatomic Markov decision processes with total reward criteria. *Journal of Mathematical Analysis and Applications*, 247(1):45–66, 2000.

[55] C. H. Fine and E. L. Porteus. Dynamic process improvement. *Operations Research*, 37(4):580–591, 1989.

[56] Centers for Disease Control and Prevention. Deaths: Preliminary data for 2001. *National Vital Statistics Reports*, 51(5):1–45, 2002. Available from http://www.cdc.gov/nchs/nvss.htm, information and data accessed on June 1, 2004.

[57] Centers for Disease Control and Prevention. United states life tables, 2000. *National Vital Statistics Reports*, 51(3):1–39, 2002. Available from http://www.cdc.gov/nchs/nvss.htm, information and data accessed on June 1, 2004.

[58] Centers for Medicare and Medicaid Services, 2003. Available from http://cms.hhs.gov/, information and data accessed on January 14, 2003.

[59] Center for Organ Recovery and Education, 2003. Available from http://www.core.org, information and data accessed on January 22, 2003.

[60] FirstGOV for seniors, 2003. Available from http://www.seniors.gov/articles/0103/healthcosts.htm, information and data accessed on January 14, 2003.

[61] S. Frederick, G. Loewenstein, and T. O'Donoghue. Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40:351–401, 2002.

[62] A. Gafni and G. W. Torrance. Risk attitude and time preference in health. *Management Science*, 30:440–451, 1984.

[63] R. J. Gallagher and E. K. Lee. Mixed integer programming optimization models for brachytherapy treatment planning. In D. R. Masys, editor, *Proceedings of the 1997 AMIA Annual Fall Symposium. A Conference of the American Medical Informatics Associations, Philadelphia, PA*, pages 278–282, 1997.

[64] K. Garber. Controversial allocation rules for liver transplants. *Nature Medicine*, 8(2):97, 2002.

[65] E. S. Gel, W. J. Hopp, and M. P. Van Oyen. Factors affecting opportunity of worksharing as a dynamic line balancing mechanism. *IIE Transactions*, 34(10):847–863, 2002.

[66] R. D. Gibbons, N. Duan, and D. Meltzer. Inequities in liver transplant allocation. *Science*, 289:549, 2000.

[67] K. D. Glazebrook. Stochastic scheduling and forwards induction. *Discrete Applied Mathematics*, 57(2-3):145–165, 1995.

[68] M. R. Gold, J. E. Siegel, B. Russell, and M. C. Weinstein. *Cost Effectiveness in Health and Medicine*. Oxford University Press, New York, NY, 1996.

[69] P. J. Green and B. W. Silverman. *Nonparametric Regression and Generalized Linear Models*. Chapman and Hall, New York, NY, 1994.

[70] J. K. Hammitt. Discounting health increments. *Journal of Health Economics*, 12:117–120, 1993.

[71] J. M. Harrison and M. I. Taksar. Instantaneous control of a brownian motion. *Mathematics of Operations Research*, 8(3):439–453, 1983.

[72] C. M. Harvey. The reasonableness of non-constant discounting. *Journal of Public Economics*, 53:31–51, 1994.

[73] Y. Hashikura, S. Kawasaki, S. Miyagawa, M. Terada, T. Ikegami, Y. Nakazawa, K. Urata, H. Chisuwa, S. Ogino, and M. Makuuchi. Recent advance in living donor liver transplantation. *World Journal of Surgery*, 26(2):243–246, 2002.

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

[75] Q. He, E. M. Jewkes, and J. Buzacott. Optimal and near-optimal inventory control policies for a make-to-order inventory-production system. *European Journal of Operational Research*, 141(1):113–132, 2002.

[76] D. P. Heyman and M. J. Sobel. *Stochastic Models*. Elsevier Science Publications, New York, NY, 1990.

[77] J. C. Hornberger and J. H. Ahn. Deciding eligibility for transplantation when a donor kidney becomes available. *Medical Decision Making*, 17:160–170, 1997.

[78] J. K. Horowitz and R. T. Carson. Discounting statistical lives. *Journal of Risk and Uncertainty*, 3:403–413, 1990.

[79] D. H. Howard. Why do transplant surgeons turn down organs?: A model of the accept/reject decision. *Journal of Health Economics*, 21(6):957–969, 2002.

[80] R. Howard. *Dynamic Programming and Markov Processes*. MIT Press, Cambridge, MA, 1960.

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

[82] M. Hughes, E. R. Carson, M. Makhlouf, C. J. Morgan, and R. Summers. Modeling a progressive care system using a color-timed Petri net. *Transactions of the Institute of Measurement and Control*, 22(3):271–283, 2000.

[83] IALOD, 2002. Available from http://www.livingdonorsonline.org, information and data accessed on February 25, 2002.

[84] MEDTAP International. The value of investment on health care, 2004. Available from http://www.medtap.com/Products/policy.cfm, information and data accessed on June 1, 2004.

[85] IOM. *Organ Procurement and Transplantation*. National Academy Press, Washington, D. C., 1999. Available from Institute of Medicine (IOM) website, http://www.iom.edu/.

[86] S. M. R. Iravani and I. Duenyas. Integrated maintenance and production control of a deteriorating production system. *IIE Transactions*, 34(5):423–435, 2002.

[87] J. S. Ivy. A manintenance model for breast cancer detection and treatment, 2002. Working paper, Available from http://www.bus.umich.edu/Academics/ Departments/OM/om/FacultyBio.asp?id=000372223/.

[88] P. S. Kamath, R. H. Wiesner, M. Malinchoc, W. Kremers, T. M. Therneau, C. L. Kosberg, G. D'Amico, E. R. Dickson, and W. R. Kim. A model to predict survival in patients with end-stage liver disease. *Hepatology*, 33(2):464–470, 2001.

[89] E. P. C. Kao and M. Queyranne. Budgeting costs of nursing in a hospital. *Management Science*, 31(5):608–621, 1985.

[90] D. A. Katz and H. G. Welch. Discounting in cost-effectiveness analysis of healthcare programmes. *PharmacoEconomics*, 3:276–285, 1993.

[91] E. B. Keeler and S. Cretin. Discounting of life-saving and other nonmonetary efects. *Management Science*, 29(3):300–306, 1983.

[92] A. J. Kleywegt, V. S. Nori, and M. W. P. Savelsbergh. The stochastic inventory routing problem with direct deliveries. *Transportation Science*, 36(1):94–118, 2002.

[93] A. L. Korotaev, P. I. Kuznetsov, A. S. Leonov, and V. E. Rodoman. Determining optimal conditions for administering an antibacterial drug. *Automation and Remote Control*, 47(1):90–96, 1986.

[94] M. Krahn and A. Gafni. Discounting in the economic evaluation of health care interventions. *Med Care*, 31:403–418, 1993.

[95] J. Kreke, A. J. Schaefer, D. C. Angus, C. L. Bryce, and M. S. Roberts. Incorporating biology into discrete event simulation models of organ allocation. In E. Yucesan, C. H. Chen, J. L. Snowden, and J. M. Charnes, editors, *WSC '02: Proceedings of the 2002 Winter Simulation Conference, San Diago, California*, 2002.

[96] E. G. Kyriakidis. Optimal control of a simple immigration-emigration process through total catastrophes. *European Journal of Operational Research*, 155(1):198–208, 2004.

[97] E. K. Lee, T. Fox, and I. Crocker. Optimization of radiosurgery treatment planning via mixed-integer programming. *Medical Physics*, 27(5):995–1004, 2000.

[98] C. Lefévre. Optimal control of a birth and death epidemic process. *Operations Research*, 29(5):971–982, 1981.

[99] L. Lladó, J. Figueras, R. Memba, X. Xiol, C. Baliellas, S. Vázquez, E. Ramos, J. Torras, A. Rafecas, J. Fabregat, C. Lama, and E. Jaurrieta. Is MELD really the definitive score for liver allocation? *Liver Transplantation*, 8(9):795–798, 2002.

[100] B. G. Lopez-Valcarel and P. B. Perez. Evaluation of alternative functional designs in an emergency department by means of simulation. *Simulation*, 63(1):20–28, 1994.

[101] G. Lyberatos and E. M. Abulesz. Analysis of periodic injections in cancer chemotherapy. *International Journal of Systems Science*, 21(8):1659–1671, 1990.

[102] P. Magni, S. Quaglini, M. Marchetti, and G. Barosi. Deciding when to intervene: A Markov decision process approach. *International Journal of Medical Informatics*, 60(3):237–253, 2000.

[103] M. Malinchoc, P. S. Kamath, F. D. Gordon, C. J. Peine, Rank J., and P. L. terBorg. A model to predict poor survival in patients undergoing transjugular intrahepatic portosystemic shunts. *Hepatology*, 31:864–871, 2000.

[104] A. D. Marcus. Extreme transplants: Doctors try radical approaches to combat organ shortage; splitting the liver in two. *Wall Street Journal*, May 6, 2003.

[105] R. B. Martin. Optimal control drug scheduling of cancer chemotherapy. *Automatica*, 28(6):1113–1123, 1992.

[106] R. B. Martin and K. L. Teo. *Optimal Control of Drug Administration in Cancer Chemotherapy*. World Scientific, River Edge, NJ, 1994.

[107] P. McMaster and B. Dousset. The improved results of liver transplantation. *Transplant International*, 5:125–128, 1992.

[108] B. J. McNeil, R. Weichselbaum, and S. G. Pauker. Fallacy of the five-year survival in lung cancer. *New England Journal of Medicine*, 299:1397–1401, 1978.

[109] MDLinx.com, 2003. Available from http://www.mdlinx.com/, information and data accessed on January 14, 2003.

[110] B. L. Miller. Dispatching from depot repair in a recoverable item inventory system: On the optimality of a heuristic rule. *Management Science*, 21(3):316–325, 1974.

[111] H. E. Miller, W. P. Pierskalla, and G. J Rath. Nurse scheduling using mathematical programming. *Operations Research*, 24(5):857–870, 1976.

[112] G. E. Monahan. A survey of partially observable Markov decision processes: Theory, models, and algorithms. *Management Science*, 28(1):1–16, 1982.

[113] T. E. Morton. On the asymptotic convergence rate of cost differences for Markovian decision processes. *Operations Research*, 19:244–248, 1971.

[114] J. Napolitano. Wisconsin senate approves tax deduction for organ donors. *The New York Times*, 2004. January 23, 2004.

[115] T. Nishizaki, S. Hiroshige, T. Ikegami, H. Uchiyama, K. Hashimoto, Y. Soejima, and M. Shimada. Living-donor liver transplantation for fulminant hepatic failure in adult patients with a left-lobe graft. *Surgery*, 131:S182–S189, 2002.

[116] Scientific Registry of Transplant Recipients (SRTR), 2003. Available from http://www.ustransplant.org, information and data accessed on January 22, 2003.

[117] Government Accounting Office, 2003. Available from http://www.gao.gov/special.pubs/organ/chapter2.pdf, information and data accessed on January 21, 2003.

[118] S. Okie. Organ exchanges push boundaries: New tactics to attract living donors raise issues of ethics and altruism. *The Washington Post*, 2001. June 9, 2001.

[119] National Institute on Aging, 2004. Available from http://www.niapublications.org/pubs/secrets-of-aging/index.htm, information and data accessed on June 1, 2004.

[120] W. P. Pierskalla and J. A. Voelker. A survey of maintenance models: The control and surveillance of deteriorating systems. *Naval Research Logistics Quarterly*, 23(3):353–388, 1976.

[121] D. K. Podolsky and K. J. Isselbacher. Chapter 298: Cirrhosis and alcoholic liver disease. In E. Braunwald, A. S. Fauci, D. L. Kasper, S. L. Hauser, D. L. Longo, and J. L. Jameson, editors, *Harrison's Principles of Internal Medicine*, pages 1752–1753, New York, NY, 2001. McGraw-Hill.

[122] E. L. Porteus. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Palo Alto, CA, 2002.

[123] F. D. Preciado-Walters, R. L. Rardin, M. Langer, and V. Thai. A coupled column generation, mixed-integer approach to optimal planning of intensity modulated radiation therapy for cancer, 2004. To appear in *Mathematical Programming*.

[124] A. A. B. Pritsker, D. L. Martin, J. S. Reust, M. A. Wagner, J. R. Wilson, M. E. Kuhl, M. D. Allen, O. P. Daily, A. M. Harper, E. B. Edwards, L. E. Bennett, J. P. Roberts, and J. F. Burdick. Organ transplantation policy evaluation. In C. Alexopoulos, K. Kang, W. R. Lilegdon, and D. Goldsman, editors, *WSC '95: Proceedings of the 1995 Winter Simulation Conference, Crystal City, VA*, pages 1314–1323, 1995.

[125] M. L. Puterman. *Markov Decision Processes*. John Wiley and Sons, New York, NY, 1994.

[126] R. Righter. A resource allocation problem in a random environment. *Operations Research*, 37(2):329–338, 1989.

[127] M. S. Roberts, D. C. Angus, C. L. Bryce, Z. Valenta, and L. Weissfeld. Survival after liver transplantation in the United States: A disease-specific analysis of the UNOS database, 2004. Forthcoming in *Liver Transplantation*.

[128] M. S. Roberts, C. L. Bryce, A. J. Schaefer, S. M. Shechter, O. Alagoz, and D. C. Angus. Optimal organ allocation policy: A biologically-based simulation model. *Medical Decision Making*, 23(6):575, 2003.

[129] M. S. Roberts and F. A. Sonnenberg. Decision modeling techniques. In F. A. Sonnenberg and G. Chapman, editors, *Decision Making in Health Care: Theory, Psychology, and Applications*, Cambridge, MA, 2000. Cambridge University Press.

[130] D. Rosenfield. Markovian deterioration with uncertain infromation. *Operations Research*, 24(1):141–155, 1976.

[131] M. W. Russo and R. S. Brown. Ethical issues in living donor liver transplantation. *Current Gastroenterology Reports*, 5:26–30, 2003.

[132] A. J. Schaefer, M. D. Bailey, S. M. Shechter, and M. S. Roberts. Medical decisions using Markov decision processes. In F. Sainfort, M. Brandeau, and W. Pierskalla, editors, *to appear in Handbook of Operations Research/Management Science Applications in Health Care*, 2004.

[133] R. L. Schmidt. A stochastic optimization model to improve production planning and R and D resource allocation in biopharmaceutical production processes. *Management Science*, 42(4):603–617, 1996.

[134] R. F. Serfozo. Monotone optimal policies for Markov decision processes. *Mathematical Programming Studies*, 6:202–215, 1976.

[135] L. S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences of the United States of America*, 39:1095–1100, 1953.

[136] S. M. Shechter, C. L. Bryce, O. Alagoz, J. E. Kreke, J. E. Stahl, A. J. Schaefer, D. C. Angus, and M. S. Roberts. A clinically based discrete event simulation of end-stage liver disease and the organ allocation process. *Medical Decision Making*, 2004. Accepted for publication.

[137] T. W. Sloan and J. G. Shanthikumar. Using in-line equipment condition and yield information for maintenance scheduling and dispatching in semiconductor wafer fabs. *IIE Transactions*, 34(2):191–209, 2002.

[138] R. Smallwood and E. J. Sondik. The optimal control of partially observable Markov processes over the infinite horizon. *Operations Research*, 21:1071–1088, 1973.

[139] E. J. Sondik. *The Optimal Control of Partially Observable Markov Processes*. PhD dissertation, Stanford University, 1971.

[140] E. J. Sondik. The optimal control of partially observable Markov processes over the infinite horizon: Discounted costs. *Operations Research*, 26:282–304, 1978.

[141] J. E. Stahl, N. Kong, S. Shechter, A. J. Schaefer, and M. S. Roberts. A methodological framework for optimally re-organizing liver transplant regions. Technical report, University of Pittsburgh, 2004. Submitted for publication.

[142] X. Su and S. Zenios. Allocation of kidneys to autonomous transplant candidates: A sequential stochastic assignment model, 2002. Conditionally Accepted by *Operations Research*. Available from http://faculty-gsb.stanford.edu/zenios/Zenios.asp.

[143] D. M. Topkis. *Supermodularity and Complementary*. Princeton University Press, Princeton, NJ, 1998.

[144] J. F. Trotter and M. J. Osgood. MELD scores of liver transplant recipients according to size of waiting list: Impact of organ allocation and patient outcomes. *Journal of American Medical Association*, 291(15):1871–1874, 2004.

[145] J. F. Trotter, N. Stolpman, M. Wachs, T. Bak, M. Kugelmas, I. Kam, and G. T. Everson. Living donor liver transplant recipients achieve relatively higher immunosuppressant blood levels than cadaveric recipients. *Liver Transplantation*, 8(3):212–218, 2002.

[146] P. A. Ubel and A. L. Caplan. Geographic favoritism in liver transplantation–unfortunate or unfair? *New England Journal of Medicine*, 339(18):1322–1325, 1998.

[147] P. A. Ubel and G. Loewenstein. Distributing scarce livers: The moral reasoning of the general public. *Social Science and Medicine*, 42(7):1049–1055, 1996.

[148] UNOS, 2003. Available from http://www.unos.org, information and data accessed on January 13, 2003.

[149] UNOS, 2004. Available from http://www.unos.org, information and data accessed on May 24, 2004.

[150] C. Valdez-Flores and R. M. Feldman. A survey of preventive maintenance models for stochastically deteriorating single-unit systems. *Naval Research Logistics*, 36:419–446, 1989.

[151] Z. Valenta. *Estimation of Survival Function for Gray's Piecewise-Cosntant Time-Varying Coefficients Model*. PhD dissertation, University of Pittsburgh, 2002.

[152] W. K. Viscusi. Discounting health effects for medical decisions. In Sloan F. A., editor, *Valuing Health Care*, Cambridge, 1995. Cambridge University Press.

[153] D. M. Warner. Scheduling nursing personnel according to nursing preference: A mathematical programming approach. *Operations Research*, 24(5):842–856, 1976.

[154] HowStuffWorks Website, 2004. Available from http://health.howstuffworks.com/organ-transplant2.htm, information and data accessed on June 14, 2004.

[155] A. S. Wein. Random yield, rework and scrap in a multistage batch manufacturing environment. *Operations Research*, 40(3):551–563, 1992.

[156] M. C. Weinstein and W. B. Stason. Foundations of cost-effectiveness analysis for health and medical practices. *New England Journal of Medicine*, 296:716–721, 1977.

[157] H. J. Weiss. The computation of optimal control limits for a queue with batch services. *Management Science*, 25(4):320–328, 1979.

[158] H. J. Weiss. Waiting time distribution in a queue with a bulk service rule. *Opsearch*, 19(1):15–24, 1981.

[159] H. J. Weiss and S. R. Pliska. Optimal policies for batch service queueing systems. *Opsearch*, 19(1):12–22, 1982.

[160] D. J. White. Dynamic programming, Markov chains, and the method of successive approximations. *Journal of Mathematical Analysis and Applications*, 6:373–376, 1963.

[161] D. J. White. *Dynamic Programming*. Holden-Day, San Francisco, CA, 1969.

[162] R. H. Wiesner, S. V. McDiarmid, P. S. Kamath, E. B. Edwards, M. Malinchoc, W. K. Kremers, R. A. F. Krom, and W. R. Kim. MELD and PELD: Application of survival models to liver allocation. *Liver Transplantation*, 7(7):567–580, 2001.

[163] H. Y. Yoo, A. Maheshwari, and P.J. Thuluvath. Retransplantation of liver: Primary graft nonfunction and hepatitis C virus are associated with worse outcome. *Liver Transplantation*, 9(9):897–904, 2003.

[164] S. Yoo, M. E. Kowalok, B. R. Thomadsen, and D. L. Henderson. Treatment planning for prostate brachytherapy using region of interest adjoint functions and a greedy heuristic. *Physics in Medicine and Biology*, 48(24):4077–4090, 2003.

[165] S. A. Zenios, G. M. Chertow, and L. M. Wein. Dynamic allocation of kidneys to candidates on the transplant waiting list. *Operations Research*, 48(4):549–569, 2000.

[166] S. A. Zenios, L. M. Wein, and G. M. Chertow. Evidence-based organ allocation. *The American Journal of Medicine*, 107(1):52–61, 1999.

[167] P. P. Zhang, J. Wu, D. Dean, L. Xing, J. Y. Xue, R. Maciunas, and C. Sibata. Plug pattern optimization for gamma knife radiosurgery treatment planning. *International Journal of Radiation Oncology Biology Physics*, 55(2):420–427, 2003.

[168] P. H. Zipkin. *Foundations of Inventory Management*. McGraw-Hill/Irwin, Boston, MA, 2000.