Applications of Stochastic Optimization Models in Patient Screening and Blood Inventory Management

by

Alireza Sabouri Bagh Abbas

B.Sc., Sharif University of Technology, 2007
M.Sc., North Carolina State University, 2009

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate and Postdoctoral Studies

(Business Administration)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

September 2014

Abstract

This thesis comprises three chapters with applications of the stochastic optimization models in healthcare as a central theme. The first chapter considers a patient screening problem. Patients on the kidney transplant waiting list are at higher risk for developing cardiovascular disease (CVD), which makes them ineligible for transplant. Therefore, transplant centers screen waiting patients to identify patients with severe CVD. We propose a model for finding screening strategies, with the objective of minimizing sum of the expected screening cost and the expected penalty cost associated with transplanting an organ to an ineligible patient. Our results suggest that current screening guidelines, which are only based on patients' risk for developing CVD, are significantly dominated by policies that also consider factors related to patients' waiting time.

In the second chapter, we extend our results from the first chapter to the case of inspecting a vital component which is needed at a random future time when an emergency occurs. If the component is not operational at that time, the system incurs a large penalty, which we want to avoid through inspections and replacements. We propose a model and solution algorithm for finding an inspection policy that minimizes the infinite horizon discounted expected penalty, replacement, and inspection costs. We also discuss other structural properties of the solution, as well as insights based on numerical results.

In the third chapter, we consider inventory decisions regarding issuing blood in a hospital. This research is motivated by recent findings in medicine that the age of transfused blood can affect health outcomes, with older blood contributing to more complications. Current practice at hospital blood banks is to issue blood in order from oldest to youngest inventory, so as to minimize shortage. However, the conflicting objective of reducing the age of blood transfused requires an issuing policy that also depends on the inventory of units of different ages. We propose a model that balances the trade-off between the average age of blood transfused and the shortage rate. Our numerical results suggest we can significantly reduce the age of transfused blood with a relatively small increase in the shortage rate.

Preface

Modified versions of Chapter 2 and 3 have been submitted for publication, but they have not been accepted for publication yet. These two papers are co-authored with Dr. Tim Huh and Dr. Steven Shechter. They were involved in the early stages of problem formulation, provided feedback during the course of both research projects, and contributed to manuscript edits. I was responsible for writing the majority of these papers, developing and implementing all the models, and preparing all the numerical results.

Table of Contents

Appendices

List of Tables

List of Figures

Acknowledgements

My sincere thanks go to my supervisors, Dr. Tim Huh and Dr. Steven Shechter at the Sauder School of Business for their invaluable help and support during my PhD studies. They helped me through every step of this research and without their guidance and contribution, this dissertation certainly would not have been the same.

I am thankful to the two other members of my committee, Dr. Maurice Queyranne and Dr. Nick Bansback for their insightful comments and suggestions. I also had the opportunity to learn from Dr. Queyranne over the course of four PhD courses and was greatly inspired by his knowledge and teaching style.

I would also like to thank my wife Behnaz and my parents for their support and encouragement throughout these years which made all this possible.

To Behnaz and Sophia

Chapter 1

Introduction

Health care has become a major application for mathematical models and the analytical tools developed in the field of Operations Research (OR) over the past decade. Health care is an area where the impact of providing better solutions is most tangible. Even minor improvements at the operational level have significant potential savings due to the high costs of the health care system. Then, these savings can compensate for the rise in the costs, standards and the need for better health care delivery. Moreover, by better managing limited health care resources, more patients can have access to high quality care in a timely manner. For these reasons, the main theme of this dissertation is to develop analytical frameworks for three important problems faced by health care policy-makers and practitioners. In the remainder of this chapter we briefly describe and motivate each problem, discuss the objectives of our work, and outline main results of our models.

1.1 Screening Strategies for Patients on the Kidney Transplant Waiting List

Chapter 2 considers the management of patients who wait for several years on the kidney transplant waiting list. Approximately 90,000 patients with end-stage renal disease await a kidney transplant in the U.S., with a median waiting time of between 2 and 5 years depending on blood type (OPTN [51]). During this time, waiting patients are at significant risk of developing cardiovascular disease (CVD), which is often described as a silent disease (McCullough [38]). Therefore, transplant centers screen waiting patients at various intervals to identify such patients and decrease the chances of operative mortality or poor post-transplant outcomes (Humar et al. [26]). Although the importance of regular cardiovascular screening of these patients is well recognized, there is no consensus on which patients should be screened and at what intervals (Gaston et al. [23]). Some centers screen all patients annually, some biannually, and some screen different patients at different frequencies according to risk level (Danovitch et al. [17]).

Since screening patients at the time of kidney arrival is infeasible (due to the limited time available), transplant centers screen patients periodically. Collaborating with a kidney transplant surgeon and transplant coordinators at the British Columbia Transplant Society, we develop a model to decide how often patients with different risk profiles and positions on the waiting list should be screened in order to minimize the expected screening cost and the cost of offering a kidney to a patient with unknown CVD. Using a simulation model, we also show that current screening guidelines, which are based only on patients' risk for developing CVD, are significantly dominated by policies that also consider factors related to patients' waiting time. We also establish that the screening intervals should be decreasing over time. This is an intuitive property of the optimal screening policy; nonetheless, it is not considered by the current screening guidelines.

To our knowledge, this is the first study that uses OR techniques to address these types of screening problems. The problem of patient readiness in the context of transplant waiting list management has its own unique complexities that make it distinct from other disease screening problems in the OR literature. The potential benefits for the patients include decreasing the number of cardiac events after renal transplant due to unidentified CVD, better use of scarce donated organs, and more efficient use of system resources for performing patient screenings.

1.2 Inspecting a Vital Component Needed upon Emergency

In Chapter 3, we extend some of our results from Chapter 2 to the case of inspecting a vital component which is needed at a random future time when an emergency occurs. If the component is not operational at that time, the system incurs a large penalty, which we want to avoid through inspections and replacements. While we describe this problem as finding inspection policies for a general component needed at the time of emergency, a health care related application for this problem is the inspection of defibrillator units placed in public buildings. These unit might fail and their failure are hidden and revealed only by inspection or attempted usage. However, if the component is not operational at the time of the emergency, the system incurs a large cost. Therefore, regular inspection and replacement of these units are important.

We propose a model and solution algorithm for finding an inspection policy that minimizes the infinite horizon discounted expected penalty, replacement, and inspection costs for these units. We also prove several structural properties of the optimal solution, and use them to develop an efficient solution algorithm. Finally, we discuss important managerial insights based on our numerical experiments. In particular, we show a property of the optimal inspection policies that contrasts with the existing literature and intuition.

1.3 Issuing Policies for Hospital Blood Inventory

The third health care problem that we study in this thesis deals with issuing policies of blood units in the hospital blood bank. Our research is motivated by the recent clinical evidence that suggests using older blood for transfusion results in higher risk for complications (Koch et al. [33], Wang et al. [76], Zallen et al. [79]). Whereas all the traditional studies of issuing policies and the current practice only consider minimizing shortages and outdates as part of their objective (Pierskalla and Roach [58]), our model also considers the age of blood transfused as one of the performance metrics that we want to optimize. The conflicting objective of reducing the age of blood transfused requires a more complicated issuance policy that also depends on the inventory of units of different ages.

Since consideration of the age of transfused blood in designing issuing policies is fairly recent, there are only a few studies in the OR literature that explore the trade-offs between the average age of blood transfused and the shortages/wastage. For instance, Atkinson et al. [7] proposed a simple class of issuing policies based on a single age threshold. They used simulation to show how changing the threshold affects the average age of blood transfused and the shortage proportion. Abouee-Mehrizi et al. [1] provided a stylized queueing model to find the distribution of the age of transfused blood under the above threshold policy. However, none of these papers characterize the optimal issuing policy or provide a framework for finding good issuing policies that consider the inventory of different ages on each day.

We formulate this problem as an infinite horizon dynamic program. Since the dynamic programming formulation of this problem suffers from the curse of dimensionality due to the large state and action spaces, we solve the problem using approximate dynamic programming (ADP) methods. Our numerical results, based on data from a large hospital in British Columbia, suggest we can significantly reduce the age of transfused blood with a relatively small increase in the shortage rate.

Our results provide evidence for policy makers that they can improve the health outcomes from blood transfusions by decreasing the average age of blood transfused without sacrificing the availability of the blood supply. Furthermore, we showed that our state-dependent ADP-based issuing policy outperforms the static policies previously proposed in the literature. We also introduce a simple class of issuing policies inspired by our results, that not only performs well, but is also easy to implement and follow in practice.

Chapter 2

Screening Strategies for Patients on the Kidney Transplant Waiting List

2.1 Introduction

End-stage renal disease (ESRD) is the complete or near complete failure of the kidney's function and marks the final stage of chronic kidney disease (CKD). According to the US Renal Data System annual report, approximately 400,000 patients in the United States had ESRD in 2009 (USRDS [74]). There are two treatment options for ESRD, dialysis and kidney transplantation, with the latter being the treatment of choice due to its lower costs and better outcomes (Port et al. [60] and Schnuelle et al. [66]). However, a shortage of living and deceased kidney donations leads to long transplant waiting lists while patients undergo dialysis. Approximately 90,000 patients in the United States await a kidney transplant, with a median wait time between 2 and 5 years depending on blood type (OPTN [51]).

All ESRD patients must pass a comprehensive evaluation process before they can be accepted to the waiting list. This includes tests for infection, cancer, and heart disease, among others. While accepted patients may be initially free of these comorbidities, during the long wait for a transplant, they may develop serious health conditions. Some of these would make the patients ineligible for transplant, since performing a major surgery on them can put them at an increased risk of mortality or serious post-transplant complications (Humar et al. [26]). Transplant centers try to avoid these outcomes by periodically re-evaluating waiting patients and temporarily putting patients they identify with such serious conditions on "inactive" (or "hold") status (Danovitch et al. [18]).

Among the different health conditions that waiting patients might develop, transplant centers are particularly concerned about cardiovascular disease (CVD), as patients typically are not aware they have developed it until they undergo a screening test (Parfrey and Foley [52]). Studies have shown that patients with CKD and patients on dialysis are at significantly increased risk of developing CVD (McCullough [38], Parfrey and Foley [52]). Also, several studies have noted that many complications after kidney transplant are associated with CVD (Humar et al. [26], Ojo et al. [50]). For instance, Ojo et al. [50] reports that in a population-based survival analysis of United States patients with ESRD transplanted between 1988 and 1997 (a total population of 86,502 adult renal transplant recipients), almost half of the deaths with a functioning graft in the first 30 days after transplant were due to cardiovascular complications (the total number of deaths with functioning graft in this period was 7,040 deaths). This leads to a firm consensus among transplant centers regarding the importance of follow-up cardiac evaluations (Gaston et al. [23]). However, while there are recommended guidelines for CVD screening policies for patients on the waiting list, they do not appear evidence-based. Furthermore, transplant centers vary considerably in how frequently they screen patients, as they use a combination of expert judgment, resource considerations, and general guidelines to come up with their own policies for screening patients (Danovitch et al. [18]).

There are three main issues that make screening decisions challenging. First, when a deceased donor kidney becomes available, it must be transplanted in a short period of time due to the fact that the prolonged time outside the body (cold ischemia time) reduces its functionality in the recipient (Salahudeen et al. [63]). Therefore, transplant centers do not have time to start screening patients at the time a kidney is offered for donation. The second issue is that one cannot know in advance with certainty to which patient the next available kidney will be offered. A kidney cannot be offered to patients whose body may develop an immune response to the donated kidney or whose blood type is not compatible with that of the donor, both of which are determined only upon an organ becoming available for donation. The last issue that complicates the screening process is that screening resources are limited and costly. Obviously, by screening all the patients in short intervals, transplant centers can maintain updated information about each patient. However, one must weigh the benefits of the extra screenings against the costs and patient inconveniences of obtaining them. These issues raise the important questions of which patients should be screened and how the screening intervals should change as patients continue waiting. Our primary contribution in this chapter is to recommend an analytical approach to answer these questions.

The rest of this chapter is organized as follows. First we review relevant literature in Section 2.2 and further describe our contributions to the literature. In Section 2.3, we state our assumptions and present our formulation of the problem. We derive several structural properties of the optimal solution in Section 2.4, and we exploit these properties to develop a solution algorithm that finds the optimal solution. In Section 2.5, we apply our model in the context of screening patients on the kidney transplant waiting list. Finally, we discuss concluding remarks and future work in Section 2.6.

2.2 Literature Review

In this section, we review the relevant literature in two parts (clinical literature and operations management literature), and we elaborate on our contribution in light of the existing literature.

2.2.1 Clinical literature

The issue of maintaining and monitoring the kidney transplant waiting list has been a challenge for the transplant community due to the persistent rise in the number of patients on the waiting list. Danovitch et al. [18] discussed this issue and summarized the results of a survey from 192 transplant centers. The focus of this survey was on screening for CVD and indicated that there is considerable variability among centers regarding the frequency and modality of the tests. They concluded that existing screening practice is not based on specific evidence and provided general recommendations for transplant programs. Gaston et al. [23] summarized the issues and recommendations regarding the kidney transplant waiting list addressed at a national meeting of the transplant community in 2002. It is mentioned in their report that “there is a widespread agreement among transplant programs that repeated cardiovascular surveillance is required for many patients awaiting a cadaver kidney transplant, with more intense monitoring for high-risk patients. There is no firm consensus, however, as to who should be tested, at what interval and with what modality.” For cardiac testing, the National Kidney Foundation [48] provided a set of recommendations that reflect the current screening policy in the U.S. (summarized in Table 2.1). In British Columbia, Canada (the geographic region we model in our numerical experiments of Section 2.5), on the other hand, the current policy is even simpler: screen “low-risk” patients every two years and “high-risk” patients annually, where high-risk is defined as a patient who is diabetic, age 50 or older, or who has a history of CVD treatment. Gill et al. [24], reporting on a study of British Columbia screening policy, concluded that

Table 2.1: Recommendations for cardiac surveillance of waitlisted patients.

Category	Screening interval
Low-risk	every 3 years
Non-diabetic high risk	biannual
diabetic and/or with history of CVD treatment	annual

further studies to obtain optimal cardiac screening policies is necessary given the high cost of performing such tests and the uncertainty in the use of the current policies highlighted in their study.

2.2.2 OR/MS literature

Our work also relates to three bodies of work in the OR/MS literature: organ transplantation, disease screening, and machine maintenance.

a) Transplant policies: The OR/MS techniques have proven helpful in developing policies for many healthcare problems, including problems related to organ transplantation. The research focus of the OM community regarding the latter has been primarily in designing better allocation policies from a societal perspective (e.g., Bertsimas et al. [14], Su and Zenios [71, 72], Zenios [80], Zenios et al. [81]) or in improving accept/reject decisions of organ offers from a patient perspective (e.g., Alagoz et al. [3, 4, 5], Sandikci et al. [64]). For instance, Zenios et al. [81] proposed a fluid model to find kidney allocation policies that optimize both clinical efficiency and equity, Su and Zenios [71] introduced a model for organ allocation that accounts for the patients' choice in accepting the organ, and Alagoz et al. [5] (in the context of liver transplant waiting list) provided a model to help patients and their physicians decide whether to accept an organ offer of given quality.

b) Disease screening policies: Several OR/MS papers have also examined optimal screening policies for different diseases. For example, models of breast cancer screening have been studied in Ayer et al. [9], Kirch and Klein [32], Maillart et al. [35], Shwartz [70]. Maillart et al. [35], for example, used a partially observable Markov chain formulation to examine the value of dynamic screening policies in which the length of screening interval can be a function of patient age. Ayer et al. [9] provided a partially observable Markov decision process (POMDP) model to determine patient-specific mammography screening times.

c) Machine maintenace policies: Our work is also related to the literature on machine maintenance polices and inspection policies in particular, as there are natural analogies between deciding when to inspect deteriorating equipment over time and when to screen patients for disease. For instance, Zhang et al. [82], building upon this literature, developed a model for finding post-operative surveillance schedules for patients who have undergone vascular surgery. We review several relevant machine reliability papers here and refer the reader to the following articles for an exhaustive survey of this literature (Chelbi and Ait-Kadi [15], Jardine and Buzacott [27], McCall [37], Nakagawa [47], Pierskalla and Voelker [59], Sherif and Smith [69], Valdez-Flores and Feldman [75], Wang [77]).

Barlow et al. [11] examined the trade-off between inspection costs and a penalty cost, where the latter is proportional to the duration of time that the system runs while a failure is undetected. They developed structural properties of the optimal solution, as well as a method to find the optimal inspection times for a broad class of failure densities. Their framework is relevant to the case of a manufacturing system in continuous operations, where failure results in the production of defective items, or in the healthcare context where early detection of disease can result in a more effective treatment. Many extensions of the Barlow et al. [11] have been considered in the literature. For instance, Luss and Kander [34] extended their model to the case where the inspection time is not negligible, and Parmigiani [55] considered the case where inspections are prone to error. Also, Munford and Shahani [39] provided a method for finding a nearly optimal inspection policy for the model studied by Barlow et al. [11], which is easier to compute and performs well. Sengupta [68] studied the case where failure becomes evident after some random time due to the evident loss of product quality or, in the patient context, due to the eventual development of symptoms.

Whereas the above papers consider a continuous running cost for the duration of time a failure goes undetected, other reliability papers consider costs that occur at discrete times during a failed or deteriorated state. For example, Maillart and Pollock [36] considered a two-phase system in which a component moves from “new” to “worn” after a random amount of time, and from “worn” to “failed” after another random amount of time. Costs accrue for each inspection, preventive maintenance action (taken when an inspection finds a component in the worn state), and reactive maintenance action (taken upon component failure), and the objective is to minimize the total expected cost over a finite time horizon. Parmigiani [54] proposed a model for finding optimal inspection policies for a stand-by component (e.g., power generator), which will be needed at a random future time (e.g., time of an emergency). In this case, there is no penalty for a delayed detection of failure, as long as it is detected before the standy-by component is needed (at which time it can be replaced). However, if the stand-by component is in a failed state at the time it is needed, the consequences are severe. The objective in this setting is to minimize the expected number of inspections, subject to a maximum probability of finding the component is failed at the time it is needed.

2.2.3 Contributions to the literature

Our work contributes to the clinical literature by proposing the first analytical framework for evaluating the performance of different screening guidelines. By integrating our analytical results at a patient-level perspective (Section 2.4) with a detailed simulation of the entire transplant waitlist (Section 2.5), we show that our model-based results can significantly outperform existing policies and other simple classes of policies. We believe this provides a foundation for developing more effective, evidence-based guidelines.

To our knowledge, our model is the first work in the OR/MS literature that addresses transplant waiting list management questions before an organ becomes available rather than after (i.e., allocation and acceptance/rejection decisions discussed above). Our problem considers how to best screen the waiting list for the *possibility* of a kidney offer, which leads to different objectives and methodologies for addressing them. Our screening problem is also different from previous disease screening studies in the literature because the goal in those is early detection and treatment of the disease while considering the cost of performing the tests. In contrast, we consider the objective of a transplant center, which is to identify adverse health conditions before a patient is called upon for a transplant, while also considering costly screening resources. Therefore, in this study, we seek to minimize the total cost from screenings as well as the cost associated with complications due to transplants to patients with unknown CVD.

Our analytical model and results have important similarities and differences to some of the reliability papers reviewed above. In particular, our two main results (Theorem 1, decreasing inspection intervals, and Theorem 2, an algorithm for obtaining the optimal policy) were derived by adapting similar results found in Theorems 5 and 6 of Barlow et al. [11]. Note that other extensions of the Barlow paper also used similar techniques to prove the optimality of decreasing inspection intervals and to provide a corresponding solution algorithm (e.g., see Theorems 2 and 3 of Sengupta [68] and Theorems 1 and 2 of Parmigiani [55]).

Our contributions to the reliability literature are in 1) providing a different extension to the original inspection problem of Barlow et al. [11], which is motivated by a health care application, and 2) filling a gap in the results that motivated previous solution algorithms in the literature. Regarding the first part, our problem requires the consideration of three interacting probability distributions: the time until the patient develops CVD, the time until the patient will be offered a kidney and undergo transplantation, and the time until the patient will die. The time until CVD development is analogous to the time until component failure (Barlow et al. [11]), and the time until calling upon the patient for a transplantation is analogous to the time until a stand-by component is needed (Parmigiani [54]). To our knowledge, there has been no analogous consideration of our patient death distribution. This might occur in a reliability context, for example, if there were some random time until the component (and any identical replacements) became obsolete.

Regarding our second contribution to the reliability literature, we note that previous results (e.g., Theorem 6 of Barlow et al. [11]) are incomplete in that they do not prove "if and only if" statements. For example, in our Theorem 2 part A of Section 2.4, we have the statement "there exists some m such that $x_j > x_{j-1} > 0$ for all $j \geq m$ if and only if $t_1 > t_1^*$." The analogous statement in the solution algorithm of Theorem 6 of Barlow et al. [11] has the "if" part of this but does not establish the "only if" part. The same holds for our Theorem 2, part B. While it is not difficult to establish the forward directions (nor would it be for the analogous result of Barlow et al. [11]), it has significant practical importance for creating the binary search solution algorithms. One will not know in general whether or not the initial choice of t_1 is greater or less than t_1^*, but rather will conclude this based on which of the two contradictions arises from the sequence of x_j that the choice of t_1 produces.

2.3 Problem Formulation and Assumptions

We first develop a model for the problem of screening an arbitrary patient on the waiting list. In Section 2.5, we discuss how this model, along with a discrete event simulation (DES) model, can be used to help a transplant center decide on screening schedules for all waiting patients. Provided the patient survives long enough on the waiting list and is still regarded as eligible for transplant, a kidney will eventually be offered to this patient. Let T be the random variable representing this remaining waiting time. In the meantime, a patient might develop CVD without knowing it, and we let S represent the random variable for the time until this would happen, provided the patient survives that long. Furthermore, the patient might also die of any reason (CVD-related or not) randomly at time U. For the sake of tractability, we assume T, S and U are independent random variables and have density functions $g(t), f(s)$ and $h(u)$, respectively. Later, we relax this assumption in our simulation model by considering dependence between the CVD development time and the death time.

If the decision maker (the transplant center) offers a kidney to the patient after developing CVD ($S < T < U$) but before screening identifies it, a penalty cost c_p is incurred due to the increased risk of complications during and after the transplant. In other words, c_p is the expected treatment cost for complications that result from offering a kidney to a patient with CVD. We let c_s represent the cost incurred each time the transplant center screens a patient. We further discuss quantifying cost parameters in Section 2.5.5. We assume c_s is significantly less than c_p. Our goal is to find a *screening policy* that minimizes the expected total cost, where a screening policy is defined by an increasing sequence of screening times $\mathcal{T} = \{t_j\}_{j=1}^{+\infty}$, where t_j refers to the scheduled time for the j^{th} screening given that the patient is still on the waiting list without any CVD detected.

We assume that the support of T, S and U is $[0, +\infty)$ and $f(s)$, $g(t)$, and $h(u)$ are positive for all positive s, t, and u. We assume $h(u)$ has the increasing failure rate (IFR) property, which is a natural property of aging. Moreover, $f(s)$ and $g(t)$ are continuous (and differentiable) *Polya Frequency* functions of order 2 (PF_2) density functions. Note that PF_2 (or log-concave) density functions include a wide range of probability densities, including the Exponential, Normal and Logistic distributions, as well as a large subset of the Gamma and Weibull family. These distributions also have the IFR property, which makes them suitable for modeling the CVD time of aging patients. Similar logic applies to the kidney offer time density since patients with longer waiting times are more likely to receive a kidney offer. Finally, we assume that screenings are error-free and take negligible time to perform. We relax these assumptions in our simulation model of Section 2.5.

2.3.1 The expected cost of a screening policy

The *expected total cost* ($\mathcal{C}$) of a screening policy is the sum of two costs: the expected penalty cost ($\mathcal{P}$) and the expected screening cost ($\mathcal{S}$). First, consider an arbitrary screening schedule, indicated by $\mathcal{T}$. To derive $\mathcal{P}$, we just need to consider the probability of incurring the penalty cost between consecutive screening times. The penalty cost c_p is incurred in the interval

(t_{j-1}, t_j) if a kidney offer is made in this interval (say at time t), the patient has developed CVD in the interval (t_{j-1}, t) before the kidney arrival, and the patient is still alive at time t:

$$\begin{aligned}\mathcal{P}(\mathcal{T}) &= c_p \sum_{j=1}^{+\infty} \mathbb{P}(t_{j-1} < S < T < t_j, U > T) \\ &= c_p \sum_{j=1}^{+\infty} \int_{t_{j-1}}^{t_j} g(t)\big(F(t) - F(t_{j-1})\big)\bar{H}(t)dt, \qquad (2.1)\end{aligned}$$

where $t_0 = 0$. In deriving (2.1), we implicitly assume $\lim\limits_{j\to\infty} t_j = \infty$, which means that screenings are scheduled as long as the patient is on the waiting list.

Now consider the expected screening cost, $\mathcal{S}$. We perform the j^{th} screening if the patient has not developed CVD by the time of $(j-1)^{th}$ screening and the kidney arrival and death have not occurred by the time of j^{th} screening:

$$\begin{aligned}\mathcal{S}(\mathcal{T}) &= c_s \sum_{j=1}^{+\infty} \mathbb{P}(T > t_j, S > t_{j-1}, U > t_j) \\ &= c_s \sum_{j=1}^{+\infty} \bar{F}(t_{j-1})\bar{G}(t_j)\bar{H}(t_j). \qquad (2.2)\end{aligned}$$

By combining (2.1) and (2.2) we obtain the expected total cost of a screening policy $\mathcal{T}$ as $\mathcal{C}(\mathcal{T}) = \mathcal{P}(\mathcal{T}) + \mathcal{S}(\mathcal{T})$. In the next section, we use this form of $\mathcal{C}(\mathcal{T})$ to establish structured optimal policies and an algorithm for obtaining the optimal policy.

2.4 Structural Results

In this section and Appendices A and B, we provide several properties of the optimal screening policy and use these properties to develop an algorithm for obtaining the optimal screening policy. Again, the perspective is that of the optimal screening policy for an arbitrary patient on the waiting list; we embed this single-patient optimal algorithm into a heuristic for screening the entire waiting list in Section 2.5. To this aim, we consider the first order optimality condition. The expected total cost $\mathcal{C}(\mathcal{T})$ is differentiable with respect to each t_j since $F(s), G(t)$ and $H(u)$ are differentiable functions (differentiability follows from the continuity of $f(s), g(t)$ and $h(u)$, respectively). Appendix A provides supporting Lemmas for the two theorems stated in this section, and Appendix B provides the proofs of these theorems.

First, we need to impose the following technical assumption on the parameters of our model. While this assumption is needed for our analysis, it holds for many instances of the problem that we solve in Section 2.5. In fact, Lemma 4 in Appendix A shows that for the case where g and h are Weibull densities, this assumption reduces to a condition only on the shape parameters of such densities. In Section 2.5, we also discuss a method to obtain nearly optimal screening policies for the cases where this assumption does not hold.

Assumption 1 *The function $\Lambda(t) := c_p - c_s\big(1 + \frac{h(t)\bar{G}(t)}{g(t)\bar{H}(t)}\big)$ satisfies the following two properties:*

(a) $\Lambda(t)$ is increasing in t.

(b) There exists some $\hat{\psi}$ such that $\Lambda(t)\bar{H}(t)g(t)$ is strictly decreasing in t for all $t > \hat{\psi}$.

Lemma 3 in Appendix A establishes that the optimal screening policy must satisfy the first order necessary condition and cannot be a boundary solution (i.e., $0 < t_1 < t_2 < t_3 < \ldots$). This Lemma also provides a tool for characterizing the optimal screening policy. To be more specific, using the necessary condition we are able to reduce the dimension of the decision variable space from $+\infty$ to 1. The necessary condition (A.3) in Appendix A produces a relationship among the three subsequent values of t_j^*'s, i.e., $\{t_{j-1}^*, t_j^*, t_{j+1}^*\}$. Since $t_0 = 0$, we can use (A.3) to find recursively all the screening times by just determining the optimal first screening time t_1^*. Therefore, for screening policies that satisfy the necessary condition (A.3), the screening times are implicitly functions of t_1^*.

For ease of representation, we define the length of the j^{th} screening interval as $x_j := t_j - t_{j-1}$. Next we show that the length of the optimal screening interval x_j^* is always decreasing in j. In other words, we screen the patient more frequently as time passes. This property is consistent with the increasing failure rate property of the disease, as well as the fact that patients are more likely to receive a kidney offer as they move up the waiting list.

Theorem 1 *In the optimal screening policy $\mathcal{T}^* = \big\{t_j^*\big\}_{j=1}^{\infty}$, the lengths of the screening intervals $\big\{x_j^*\big\}_{j=1}^{\infty}$ are decreasing in j.*

The decreasing-interval property proved in Theorem 1 is the most important property of the optimal screening policy, and yet it is ignored by all of the constant-interval screening guidelines used in practice (see Section 2.2.1). In our numerical experiments of Section 2.5, we show that constant-interval policies are significantly outperformed by our decreasing-interval screening policies.

Our next step is to obtain the optimal screening policy. As we discussed, the optimal screening times can be obtained recursively using (A.3) after obtaining the optimal first screening time t_1^*. The following result motivates our solution algorithm by providing us with a way of identifying a wrong choice for the first screening time t_1.

Theorem 2 *Let $\{x_j\}_{j=1}^{\infty}$ be the sequence of the lengths of the screening intervals generated recursively from (A.3) after choosing a first screening time, $t_1 > 0$. Let t_1^* be an optimal value of the first screening time t_1. Then, the following holds:*

(A) There exists some m such that $x_j > x_{j-1} > 0$ for all $j \geq m$ if and only if $t_1 > t_1^*$.

(B) There exists some m such that $x_j < 0$ for all $j \geq m$ if and only if $t_1 < t_1^*$.

Theorem 2 establishes that if a non-optimal first screening time t_1 is chosen and being used to generate a sequence of screening time recursively using (A.3), it would result in some "contradiction" (we say "contradiction" since $x_j > x_{j-1} > 0$ is inconsistent with Theorem 1, and $x_j < 0$ is inconsistent with the constraint $t_1 \leq t_2 \leq \ldots$). It also states that if the contradiction is of the kind "$x_j < 0$ for some j", then the current choice of t_1 is less than the optimal value t_1^*, and if the contradiction is of the kind "$x_j > x_{j-1}$ for some j", then the current choice of t_1 is greater than the optimal first screening time t_1^*. These conclusions also imply that the optimal screening policy is unique. To obtain the unique optimal first screening time t_1^*, a binary search algorithm can be used since based on the guidelines of Theorem 2, we can decide whether $t_1 > t_1^*$ or $t_1 < t_1^*$.

2.5 Numerical Experiments

This section describes how the results from previous sections can be applied to the problem of screening patients on the kidney transplant waiting list. First, we describe a data-driven discrete event simulation (DES) model we developed of a kidney transplant waiting list. This serves two purposes: 1) we can use the DES to evaluate any proposed screening policy, and 2) we use the DES to estimate the remaining waiting time density g for our analytical model, which are then used in our model-based screening policy (described in Section 2.5.3).

2.5.1 Simulation of the kidney transplant waiting list

Our DES is based on the kidney transplant waiting list of British Columbia, Canada. The main components of the model are as follows:

1. Patient generation: The patient generation module creates ESRD patients according to a Poisson process with rate 304 patients per year and assigns them various characteristics such as blood type, age, and gender. These characteristic are assigned randomly based on a probability distribution obtained from historical data (see Section 2.5.2 for data sources).

2. Health progression while waiting: The simulation program assigns each patient a time at which the patient may develop CVD, as well as a time of death. We use accelerated failure time models (Anderson et al. [6], Odell et al. [49]) for modelling the time of developing CVD and the time of death. These models have common risk factors such as age, gender and being diabetic as covariates. Since these models are calibrated based on data from the general population, we use a Cox proportional hazard model to update them for ESRD patients (for instance, Foley et al. [21] and Sarnak and Levey [65] report that CVD mortality in dialysis patients is 10-30 times higher than in the general population). We remark that here the time of CVD development means the time at which the severity of the disease reaches the point that the transplant center would no longer consider the patient eligible for a transplant. In the absence of a clear definition for this event, we assume that the definition used in Anderson et al. [6] matches the practice of transplant centers.

3. Kidney generation: Similar to the patient generation module, we use British Columbia historical data to generate the time between kidney donations (average 6.1 days, exponentially distributed) and the donor's blood type.

4. Screening policy: The screening policy indicates at what times different patients should undergo cardiovascular screening so as to update their health status and eligibility for transplant. We consider imperfect screenings through their estimated sensitivity and specificity, as well as screening lead time, which denotes the time until the results of the screening becomes available. In practice, screening is a two-stage process. In the first stage, a non-invasive test (e.g., echo-cardiogram or electrocardiogram) is performed. Patients are asked to take an invasive test (coronary angiography) only if the result of the first stage tests is positive. If the overall result of the screening process is positive, then the patient is placed on hold status at the time the screening result becomes available. Also, the first stage tests are usually imperfect (i.e., have sensitivity and specificity of smaller than 1), where false positive test results increase the screening costs due to unnecessary follow-up tests, and false negative results lead to CVD cases that go undetected. Furthermore, the final outcome of the screening process becomes available after a non-negligible lead time (which we consider as exponentially distributed with a mean of one month). Note that since for each patient the simulation program has already assigned the true CVD development time, the result of each screening test can be easily determined given the sensitivity and specificity of the tests. The simulation program assigns a cost to each screening performed, where the cost of the second-stage test is only applied if the result of the first-stage test is positive.

5. Allocation policy: After a kidney becomes available, the allocation policy indicates to which patient it should be offered. We use the allocation policy currently in place in British Columbia, which is primarily driven by the waiting time and compatibility criteria, i.e., the kidney goes to the compatible patient (cross-match negative, blood type compatible patient) who has been waiting the longest. If the kidney is offered to a patient who has developed CVD that is not identified by any of previous screenings, we assign a penalty cost c_p to this transplant (see Section 2.3 for the description of c_p).

6. Outcomes/output: The main outcome of interest in our case is the total cost. Other detailed metrics such as annual number of screenings performed or the annual number of CVD cases detected are also collected by the simulation model.

2.5.2 Data

We use publicly available data from different sources. The simulation model is based on the waiting list dynamics (patient and kidney arrival rates, as well as allocation policy) that occur in British Columbia. However, the characteristics of the patients and kidneys are generated based on both British Columbia and United States data, where the latter is used if we could not find the data for British Columbia. Patients' CVD time distributions are generated using accelerated failure time models introduced in Anderson et al. [6]. For a patient's death time, we first generate a random time from a death time distribution for a general ESRD patient (who has not developed CVD yet). If the generated time is after the patient's CVD time, we update the death time by generating it from a death time distribution for a ESRD patient who also has CVD, based on the model of Anderson et al. [6]. Table 2.2 summarizes the main model components, parameter values, and data sources.

Table 2.2: Parameter values and data sources for main model components.

	Parameter	Value	Source
Patient	Arrival rate	304 ESRD patients per year	†
	Age distribution (M)	30-40 (0.176)[1] 40-50 (0.241) 50-60 (0.401) 60-70 (0.144) 70-80 (0.037)	†
	Age distribution (F)	30-40 (0.177) 40-50 (0.301) 50-60 (0.345) 60-70 (0.168) 70-80 (0.009)	†
	Gender	Probability of female: 0.45	‡
	Blood type	O (0.485) A (0.329) B (0.148) AB(0.038)	‡
	PRA score	0-10 (0.7520) 10-80 (0.1720) 80-100 (0.0760)	‡
	Diabetic	Probability of diabetes: 0.6	‡
	CVD time	Accelerated failure time models	Anderson et al. [6]
	Death time	Accelerated failure time models	Anderson et al. [6]
Kidney	Transplant rate	60 kidney transplants per year	†
	Blood type	O (0.474) A (0.372) B (0.123) AB(0.031)	‡
Screening	Sensitivity	Electrocardiogram: 0.68, Coronary Angiography: 1.00	Patterson et al. [56]
	Specificity	Electrocardiogram: 0.77, Coronary Angiography: 1.00	Patterson et al. [56]
	Mean lead time	1 month	∗
Cost	Screening cost	ECG: $330, Coronary Angiography: $1800	Patterson et al. [56]
	Penalty cost	$40,000	Patterson et al. [56]

[1] Numbers in parentheses represent probabilities.
† Canadian Organ Replacement Register Report, 2010, and British Columbia Transplant Society (BCTS) data, 2006.
‡ Organ Procurement and Transplantation Network (OPTN) data, 2011.
∗ Discussion with transplant coordinators at BCTS.

While limited public data make it difficult to perform extensive validation of our model, we were able to compare observed median waiting time with the ones reported by the British Columbia Transplant Society [12]. The median waiting time experienced in British Columbia was 63.1 months for patients transplanted in 2010 and 62.2 for patients transplanted in 2011, which are close to the median waiting time of 63.0 months observed from the simulation program.

2.5.3 Model-based screening policy

Here we discuss how we apply our analytical model of Sections 2.3 and 2.4 to develop a model-based screening policy. Note that our analytical model determines the optimal screening times for a single patient who faces a remaining waiting time (until kidney offer) density g and has CVD and death time densities f and h, respectively. In practice, a transplant center needs to develop a screening policy for all patients on the waiting list. However, it would be analytically intractable to derive a globally optimal solution for deciding all patients' screening times; the state space and decision space would quickly face the "curse of dimensionality." Instead, we propose a heuristic that first assigns each new patient a screening schedule, derived from our analytical model and solution algorithm of Section 2.4. We then dynamically update patients' screening schedules as the waiting list evolves and they move to higher ranks.

The CVD and death time densities (f and h, respectively) are functions of patient characteristics such as age, gender, and diabetes status, and are independent of the waiting list dynamics and other patients' characteristics. In contrast, the remaining waiting time density g for each patient is a function of the waiting list dynamics and depends on the characteristics of patients who have higher priority over the considered patient. This is where we use the DES model for the purpose of estimating the density g for a patient, as we can run several replications of the model to obtain waiting time samples, and then fit a distribution to the data. However, it would be computationally expensive to do this for every patient at different times during the simulation run. Instead, we create a moderate number of categories based on a few factors related to the waiting time distribution, and fit these distributions offline. These factors are a patient's rank on the waiting list, blood type, and panel reactive antibodies (PRA) score.

Rank is an important determinant of remaining waiting time, since kidneys are offered in first-in-first-out order (provided the kidney is compatible with the intended recipient). A patient's blood type is also a major factor affecting a patient's waiting time. For instance, a patient of rank 50 with blood type AB will experience a much shorter waiting time compared to a patient ranked 50 with blood type O, since blood type AB is the universal recipient and blood type O is the universal donor. The third factor, the patient's PRA score, estimates the percentage of the donor population that the patient would develop an immune response to, and thus be unable to receive the donor's kidney. Patients with lower PRA scores will wait less for a kidney offer. We consider all combinations of 30 rank (up to rank 300 in groups of 10), 4 blood type (A, B, AB, O), and 3 PRA score (0-10, 10-80, 80-100) categories, for a total of 360 waiting time distributions we fit using the DES model.

We also group patients into 20 categories based on the primary factors that affect the CVD and death time distributions f and h: 5 age (30-80 in groups of 10), 2 gender, and 2 diabetes status categories. We then use our analytical model of Sections 2.3 and 2.4 to obtain the optimal screening policy for 7,200 (360×20) different combinations of a patient's characteristics. Since we dynamically update the optimal screening policy every month, for each combination, we just store the first screening interval, t_1^* in a look-up table. If the look-up value of t_1^* is within the next month, we screen the patient. In the remainder of this chapter, we call this look-up table the "model-based" screening policy.

Having this look-up table effectively allows us to implement a dynamic screening policy. As an example, suppose that based on the current rank of a patient, the optimal policy suggests that the patient should be screened after one year. However, after six months, the transplant center might observe that the rank of the patient decreased faster than expected due to a higher donation rate in that six month period. They can then use the optimal screening policy for the updated rank, which might suggest to screen the patient much sooner than the six months that remain in the previous screening schedule for that patient.

We remark that for 14% of the 7,200 combinations, the particular set of parameters does not satisfy Assumption 1 described in Section 2.4, and therefore our solution algorithm cannot be used to find the optimal patient-specific screening policy. For these cases, we find a nearly optimal policy using a method inspired by Munford and Shahani [39]. In this method, given a probability p, we find the next screening interval in such a way that the probability of offering a kidney to a patient with unknown CVD in that interval is equal to p. After obtaining the screening intervals, the total expected cost of such a policy (as defined in Section 2.3) can be easily computed. To find the best value for p, we try different values for p and choose the one that minimizes the total expected cost. The corresponding screening policy provides an approximation to the optimal policy. When compared against the optimal policy for the cases where Assumption 1 holds, we observed that the total expected cost is only 7% less for the optimal policy on average.

Note that our analytical model in Section 2.3 only considers one parameter, c_s, for the screening cost, whereas in reality (and as our simulation model considers) there are two stages of tests, with the second stage being significantly more expensive than the first stage (see Table 2.2 for cost estimates). To deal with this difference between our analytical and simulation models, we use an implied value of c_s, which we define as the sum of the cost of performing the first non-invasive test (c_1) and the expected cost of performing the second test (ac_2), where a is the proportion of times that the second test is performed. Using the simulation program, we observe that for the current screening policy, the second test is performed 36% of the time. We use this value as our estimate of a.

2.5.4 Results and discussion

This section aims to provide insights and guidelines to policy-makers for designing more effective and efficient screening policies. Towards this aim, we perform several numerical experiments and discuss the importance of different variables in the design of improved screening policies.

We first demonstrate the value of our model-based screening policy by comparing its performance against the current guidelines of screening high-risk patients annually and low-risk patients every two years. We also compare our approach against the best among two other classes of screening policies: "risk-based" and "rank-based" policies. A risk-based policy is a generalization of the current guidelines, which assigns one fixed screening interval to high-risk patients, and another fixed screening interval to low-risk patients. We consider fixed intervals ranging from 3 months to 3 years, in increments of 3 months. For the rank-based policies, we consider a single rank threshold (50, 100, 150, 200, or 250) and assign different screening intervals to the patients on different sides of this threshold. It is worth pointing out that instead of using the model-based policy, one might consider a simulation optimization approach. However, instead of searching over an enormous decision space for good policies via common simulation optimization approaches (e.g., genetic algorithms, tabu search, etc.), we use our analytical model to find a good screening policy.

The results of our experiment are summarized in Table 2.3. We performed 100 simulation replications of each policy, and each replication simulates 1,000 kidney transplants after the warm-up period (which ends when the waiting list size matches the current waiting list size in British Columbia). The table provides the averages and the 95% confidence intervals for the following three metrics: total annual cost, percentage of bad transplants (i.e., transplants offered to patients with unknown CVD), and average annual number of screenings (both non-invasive and invasive). Note our policies focus on trying to minimize the total expected cost metric, but the other two metrics provide other outputs of interest for the transplant center.

Table 2.3: Comparison of different screening policies.

Screening policy	Total annual cost(million \$)	Percentage of bad transplants	Average annual number of screenings
Current	0.797 ($\pm$0.007)	20.9 ($\pm$0.2)	409.9 ($\pm$3.2)
Best risk-based	0.714 ($\pm$0.007)	12.6 ($\pm$0.2)	592.8 ($\pm$4.8)
Best rank-based	0.612 ($\pm$0.005)	7.2 ($\pm$0.2)	639.5 ($\pm$5.2)
Model-based	0.511 ($\pm$0.005)	8.5 ($\pm$0.2)	429.1 ($\pm$4.0)

First, we observe that the current screening policy results in 20.9% bad transplants and requires 409.9 screenings performed each year. We found that the best (lowest total cost) risk-based policy should screen high-risk patients every six months and low-risk patients every year, which is half of the intervals suggested by the current policy. Consequently, the annual number of screenings performed increases by 44.6%, but the total annual cost decreases by 10.5% due to the 8.3% decrease in the percentage of bad transplants. The best rank-based policy suggests screening the first 100 patients every 3 months, and the rest every 3 years. Compared with the best risk-based policy, this policy performs 7.9% more screenings each year, but reduces the total annual cost and the percentage of bad transplants further by 14.2% and 42.8%, respectively. This suggests the importance of considering factors highly correlated to remaining waiting time (e.g., rank on the wait list) in the design of patient screening guidelines. Our model-based policy, which explicitly considers the remaining waiting time distribution, achieves a 17% reduction in total cost relative to the best rank-based policy. It does this by by requiring significantly lower annual number of screenings with only a slightly higher percentage of bad transplants. Comparing the model-based policy back to the current policy, we obtain a 35.9% reduction in the total annual cost, by reducing the percentage of bad transplants by 59.6% while using only 5% more screenings per year. The model-based screening policy achieves this efficiency by scheduling the screenings at periods where the patient is more likely to receive an offer or have developed CVD.

2.5.5 Comparison of policies in the face of cost uncertainty:

Recall that we defined the penalty cost as the treatment cost of complications resulting from a transplant to a patient with unknown CVD. In practice, however, different complications might arise as a result of performing a transplant on a patient with severe CVD, and the treatment cost for these complications can significantly vary. Furthermore, the penalty cost can consider the opportunity cost of offering a kidney to a healthy patient instead of a patient with severe CVD. All these issues make quantifying cost parameters a difficult task in practice. Instead of focusing on the cost, Figure 2.1 provides an alternative way for managers to understand the benefit of using our model-based policy over existing guidelines. Figure 2.1 depicts trade-off curves for the two primary (non-cost-based) metrics in our study: the percentage of bad transplants and the annual number of screenings. Each point of the figure corresponds to one screening policy. For the risk-based policies, each point is obtained by choosing different fixed screening intervals for the two risk groups. Similarly, for the rank-based policies, each point is obtained by choosing a rank threshold as well as two fixed screening intervals for the two rank groups. We remark that the dominated policies in each class are not shown in Figure 2.1. The different model-based policies are obtained by changing the ratio of costs $\frac{c_p}{c_s}$.

Figure 2.1 shows that for any risk-based policy, one can find a dominating rank-based policy (i.e., a policy with a smaller percentage of bad transplants and smaller annual number of screenings performed). Furthermore, one can find a model-based policy that outperforms both risk and rank-based policies. Therefore, by using our model-based policies, one can always gain efficiencies in terms of both metrics. Our observation suggests that the insight discussed earlier about the importance of considering the waiting time in designing the screening guidelines does not change if different cost figures are used. Furthermore, the Pareto efficient curves shown in Figure 2.1 provides a powerful tool to policy-makers for developing the right screening guidelines, as they can choose a policy which balances their desired trade-off between the percentage of bad transplants and the amount of screenings necessary.

Figure 2.1: Pareto efficient policies from different policy classes. The "best" coordinate for each policy reflects the values found in Table 2.3, which are based on the baseline model parameters given in Table 2.2 .

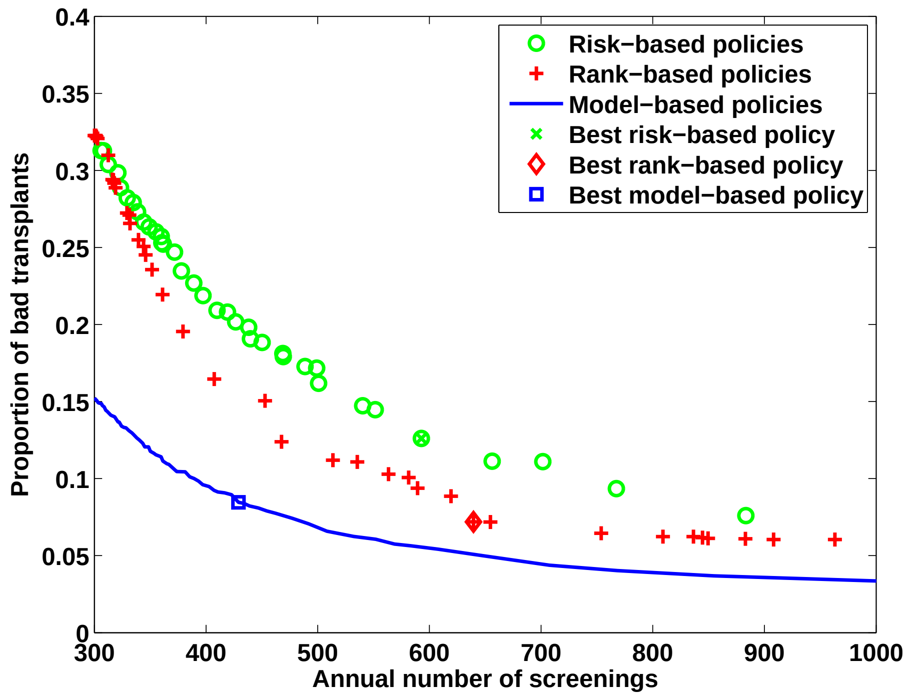

2.6 Conclusions

This chapter provides a new modeling framework for finding improved screening policies for patients on the kidney transplant waiting list. We not only provide the first evidence-based tool for designing such policies, but also highlight several important insights regarding good policies. Our model balances the trade-off between two costs: the penalty cost as a result of offering a kidney to a patient with unknown CVD and screening costs. We prove several properties of the optimal patient-specific screening policy for this model and use these properties to develop a binary search solution algorithm. In particular, we show that the lengths of the screening intervals are decreasing under reasonable assumptions. Later, with the help of a discrete event simulation model, we suggest a heuristic method that uses our single-patient model to develop dynamic screening policies within the multi-patient waiting list setting.

The current screening policies are based on simple rules which may be easier to follow in practice. However, as we discussed earlier, transplant centers view these policies as general guidelines and deviate from them by periodically reviewing each patient's information. On the other hand, the model-based policy developed by our model can be summarized in the form of a look-up table (based on age, gender, diabetes status, blood type, PRA score, and rank) that indicates how long a transplant center should wait before requesting a new screening. The screening schedule is obtained by the binary search algorithm of our analytical model, and therefore our approach may be seen as an efficient alternative to a general simulation-based optimization approach to finding screening policies. We also note that our look-up table can be easily implemented in the form of a spreadsheet, and based on our discussion with British Columbia transplant coordinators, this would significantly facilitate the screening process.

As mentioned in Section 2.3, our analytical model of Section 2.3 made some simplifying assumptions that may pose limitations: perfect testing, and zero lead time for test results. Our DES model, on the other hand, relaxed each of these assumptions and reflected the reality of the process. In Section 2.5.3, we discussed our approach for mitigating the first assumption by estimating a screening cost parameter of the model (c_s) that reflected the two stages of testing that occur in practice. The incorrect assumption of a zero lead time would only matter if a patient had developed CVD since the last screening and a kidney were offered to the patient during the time it took to receive the test results of the new screening. The probability of this happening will be negligible for patients further down the waiting list but may be worth considering for patients with shorter times until receiving a kidney offer. Despite the disconnect between aspects of our analytical model and the simulation model, our model-based screening policy performed well compared to other classes of policies.

Our numerical experiments provide several important insights regarding the effective and efficient screening policies. First, we observe that variable-interval screening policies perform better than the fixed-interval policies used in practice. Furthermore, our results suggest that the factors affecting the waiting time of the patient, such as rank and blood type, should be considered in determining the screening intervals. This contrasts sharply with current guidelines which are only based on patients' risk for developing CVD. Finally, we note that while our numerical results are based on data for the British Columbia Transplant Center, our framework can be easily applied to other regions by using appropriate data to calibrate the various model parameters.

Chapter 3

Inspecting a Vital Component Needed upon Emergency

3.1 Introduction

Consider a vital component that fails after some random amount of time and that will be needed at some other random future time when an emergency occurs. If the component is not operational at the time of the emergency, the system incurs a large penalty cost. Failures are hidden and revealed only upon inspections or upon attempted usage, at which time the component is replaced (or repaired). Furthermore, the component can be replaced preemptively after some time.

As an example, consider stand-by safety systems, such as smoke detectors and fire alarms, that are needed in emergencies. As stand-by units, their failure may go unnoticed until they are needed, which could result in very costly or even catastrophic outcomes. For instance, the National Fire Protection Association reports that 24% of fire deaths in the US between 2005 and 2009 resulted from fires in homes in which smoke alarms were present but did not operate (Ahrens [2]). Regular inspections of such systems is clearly important, though it is not obvious how frequently to perform such inspections. Maintaining emergency back-up equipment, such as power generators, is another application of this problem. For instance, during a major storm, scattered power outages can require the use of standby power systems installed at hospitals, commercial businesses, and government buildings. If the backup generators are not working at these times, there may be major consequences.

There are several important questions in designing a cost-effective maintenance plan for these systems, which arise from the trade-offs among different decisions. First, there is a trade-off between performing frequent replacements, accumulating high replacement costs, versus infrequent replacements, with higher risk of failure happening prior to the attempted use. In other words, the first decision is determining the right time to replace the component (i.e., the cycle length). This preemptive replacement can be delayed by scheduling inspections to detect the component failure. However, as the component ages and is more likely to fail, inspections must be scheduled more frequently, and their increasing cost can make delaying the replacement no longer economical. Therefore, one should decide on the right number of inspections performed in each cycle before replacing the component. Finally, one should decide how to space the inspections within a cycle so as to effectively reduce the likelihood of a catastrophic outcome (i.e., attempting to use a failed component). We answer all these questions and provide several managerial insights for designing the optimal policy.

The rest of this chapter is organized as follows. First we review the relevant literature and indicate the contribution of our work in Section 3.2. In Section 3.3, we state our assumptions and present our formulation of the problem. We derive several structural properties of the optimal solution in Section 3.4, and we exploit these properties to develop a solution algorithm. In Section 3.5, we present and discuss numerical examples. Finally, we discuss managerial insights and future work in Section 3.6. Proofs of main results appear in the appendix.

3.2 Literature Review

Optimal inspection policies have such a rich history in operations management that several articles have surveyed this literature (Chelbi and Ait-Kadi [15], Jardine and Buzacott [27], McCall [37], Nakagawa [47], Pierskalla and Voelker [59], Sherif and Smith [69], Valdez-Flores and Feldman [75], Wang [77]). Of particular relevance to our setting is a seminal paper by Barlow et al. [11] on optimal inspection policies and various extensions of the ideas therein. They considered an inspection problem in which they assumed that (a) system failure is known only through inspection, (b) inspections do not degrade the system and take negligible time, (c) a cost is associated with each inspection (d) a cost is associated with each unit of time that the failure is undetected, and (e) the problem ends upon detection of the failure. Note that assumption (d) makes this model suitable for a system in continuous operation (such as a production system), where system failure may result in producing defected items proportional to the duration of the time that the system has failed. Barlow et al. [11] developed structural properties of the optimal solution, as well as a method to find the optimal inspection times for a broad class of failure densities. Barlow and Proschan [10] relaxed assumption (e), and extended their analysis to the case where the component is replaced upon detection of failure. They considered the objective of expected cost per unit of time over an infinite horizon.

Many other extensions to Barlow et al. [11] have been considered in the literature. Luss and Kander [34] extended their model to the case where the inspection time is not negligible. Sengupta [68] considered a system for which failure becomes evident at some random time after failure. For example, in a manufacturing process, the operator eventually would detect the machine failure because of the deterioration in product quality. Parmigiani [55] considered a system where the inspections are fallible and take non-negligible time.

While the papers mentioned above assign a penalty cost for each unit of time that the failure is undetected, in our problem, a fixed penalty cost is incurred only if one attempts to use a failed component. This relates to literature examining inspection policies for stand-by systems. For example, Nakagawa [46] considered inspection policies for a stand-by unit that replaces the main unit if it fails. He assumed the inspection times are equally spaced and the problem ends when the main unit fails. With these assumptions, he considered the problem of finding the inspection interval that minimizes the expected cost, and found sufficient conditions for the existence of the optimal inspection interval. However, he did not provide any algorithm to obtain one. Whereas Nakagawa [46] studied the problem of determining the optimal inspection interval within the class of policies of equally spaced intervals, Thomas et al. [73] modeled the problem as a discrete Markov decision process to find the inspection/repair times that maximizes the expected time until a catastrophe occurs (i.e., stand-by unit is needed but it is not operational). Their model considered that inspections are fallible and the inspection and repair durations are not negligible. Sabouri et al. [62] considered the problem of screening patients on the kidney transplant waiting list and modeled it as a type of inspection model for a stand-by system.

Another stream of models developed for systems in standby or storage consider system availability (ratio of the system up time to the total of system up and down times) as the primary performance criterion instead of the average total cost. The goal in these models is to maintain a certain level of availability with minimum inspection effort. Parmigiani [54] considered the inspection of stand-by units in continuous time by minimizing the expected number of inspections under the constraint of a maximum probability that the stand-by unit is not ready upon the failure of the main unit. He used an approach similar to Barlow et al. [11], and by imposing certain assumptions on both the failure distributions, he provided an algorithm to obtain the optimal inspection times. Yeh [78] developed a model to find an optimal inspection-repair-replacement policy that maintains a certain level of availability at any time and minimizes the long-run expected cost per unit time. As suggested in Kim and Thomas [31], the models studied by Yeh [78] and several other authors did not consider the case where the need for the component changes over time. In contrast, Kim and Thomas [31] characterized the optimal repair decisions in the case where the need for the equipment varies over time according to a Markov chain. However, their model did not consider finding an optimal inspection schedule.

Our model in this chapter assumes that inspections only reveal whether the system has failed or not. While this assumption is reasonable for several applications (e.g., fire alarms), it is worth mentioning that a more recent stream of research considers situations where at the time of an inspection, the degree of equipment degradation can be determined by performing measurements. Then, preventive replacement is performed when a certain threshold of equipment degradation is reached. The focus of these conditional maintenance models (reviewed in Chelbi and Ait-Kadi [15]) is either in finding an optimal inspection policy for a given threshold or finding the optimal threshold for a given inspection schedule. As an exception, Dieulle et al. [20] proposed a model to optimize both the threshold and the inspection schedule.

Our contributions to the literature are as follows. First, instead of focusing on an availability criterion, we consider a system in which component failure is only a problem at certain times (i.e., time of emergency), and therefore it is critical that our objective function considers the interactions between the component failure distribution and the emergency time distribution. In this context, we consider the optimal design of a joint inspection and preventive maintenance policy by finding the optimal number of inspections and their timings before a preventive replacement. Second, our analysis of the case where only a finite number of inspections are scheduled reveals a property of the optimal inspection schedule which contrasts with the structure of optimal policies for models with an unlimited number of inspections (Theorem 3). We also provide insights on the impact of changes in the model parameters by solving several numerical examples. Finally, we explore the performance of constant interval (periodic) inspection policies, and provide insights on the performance loss when one is restricted to this class of policies.

3.3 Problem Formulation and Assumptions

The time between emergencies is a random variable, which we denote by T. We suppose that at the time of emergency, a vital component is needed (e.g., back-up power generator), and an attempt will be made to use it. This component deteriorates over time and has a random lifetime S (independent of T), where the probability density function of S is $f(s)$, and its cumulative distribution function is $F(s)$. If the component fails and the failure remains undetected until the emergency occurs (we refer to this event as a catastrophe), the system incurs a penalty cost, c_p, at the time of the emergency (as the component cannot be used). Failure can be detected only by inspection, at a cost of c_i per inspection. We assume that upon identifying component failure (either before or at the time of emergency), upon using the component at the time of emergency, or at a scheduled time in future, it is replaced (or repaired to become as good as new) in negligible time at a replacement cost of r (in addition to possibly c_p or c_i). We also assume that:

1. the remaining time until the next emergency is independent of the time that has passed since the last emergency (i.e., the memoryless property holds), and therefore T has an exponential probability density function with parameter λ (i.e., $g(t) = \lambda e^{-\lambda t}$).

2. $f(s)$ is a continuous (and differentiable) Polya Frequency functions of order 2 (PF_2), and the support of S is $[0, +\infty)$ and $f(s) > 0$ for all $s > 0$. Note that PF_2 density functions include a range of probability densities, including all Exponential and Normal distributions, as well as a large subset of the Gamma and Weibull family. These distributions have the increasing failure rate property, which makes them suitable lifetime distributions for components subject to deterioration over time.

3. $c_i < c_p$ and $r < c_p$; otherwise, we are not willing to perform inspections or replace the component to avoid the penalty cost. Furthermore, inspections are error-free, take negligible time to perform, and do not affect the failure rate of the component.

4. the planning horizon starts with a newly replaced component.

Our model considers that a decision maker can perform a preventive replacement if it is no longer economical to perform further inspections. This may happen if the interaction between the component failure rate and the emergency time distribution requires a very high inspection frequency to avoid the large penalty c_p. Therefore, we must identify the optimal number of inspections to perform before scheduling a preventive replacement. Toward this aim, we first formulate a subproblem for finding the optimal timing of the inspections and the replacement for a given number of inspection opportunities per cycle (the time between replacements). Then, one can solve several instances of this problem to obtain the optimal number of inspection opportunities.

Since the emergency arrival density has the memoryless property, a renewal happens each time the component is replaced. For the subproblem, we assume that there are a limited number n of inspection opportunities during each cycle, and that any time after performing n inspections, the decision maker can preventively replace the component. Our goal is to find an *inspection policy* for each cycle that minimizes the infinite horizon expected total discounted cost $\mathcal{C}$. An inspection policy is defined by an increasing sequence of inspection times $\mathcal{T}_n = \{t_j\}_{j=1}^{n+1}$ (where $t_1 \leq t_2 \leq \ldots \leq t_n \leq t_{n+1}$ and t_{n+1} is the scheduled replacement time). Note that since the problem renews itself after each cycle, the same inspection policy is optimal for each cycle. Therefore, without loss of generality, we restrict our attention to the class of policies where the same inspection policy is used for all cycles, and we set the beginning of each cycle to time 0. Then, t_j represents the time since the last renewal, and $\mathcal{C}(\mathcal{T}_n)$ represents the infinite horizon discounted expected total cost when in each cycle we use the same inspection policy $\mathcal{T}_n$. We revisit the optimal choice of n in the numerical results of Section 3.5.

3.3.1 The discounted expected total cost of an inspection policy

Each cycle ends with a replacement in one of the following four ways: (a) by identifying failure upon emergency, (b) by identifying failure upon inspection prior to an emergency, or (c) by replacing the component at the scheduled time t_{n+1}, or (d) by using the functioning component upon emergency (in which case we assume it needs to be replaced or restored to new condition at the same cost). The objective function can be written recursively based on which of these scenarios causes a component replacement. In equation (3.1), the first three summands correspond to these cases (case (c) is combined with (b)), respectively, and the last expression is the total discounted expected inspection cost in one cycle:

$$\begin{aligned}\mathcal{C}(\mathcal{T}_n) &= (c_p + r + \mathcal{C}(\mathcal{T}_n))\mathcal{A}_1(\mathcal{T}_n) + (r + \mathcal{C}(\mathcal{T}_n))\mathcal{A}_2(\mathcal{T}_n) \\ &\quad + (r + \mathcal{C}(\mathcal{T}_n))\mathcal{A}_3(\mathcal{T}_n) + c_i\mathcal{A}_4(\mathcal{T}_n). \end{aligned} \tag{3.1}$$

Now we explain and derive the different expressions in equation (3.1). The first expression, $(c_p + r + \mathcal{C}(\mathcal{T}_n))\mathcal{A}_1(\mathcal{T}_n)$, represents the discounted expected cost if a failure remains undetected until the time of an emergency (meaning $\mathcal{A}_1(\mathcal{T}_n)$ can be viewed as the discounted probability of such an event). To derive the $\mathcal{A}_1(\mathcal{T}_n)$, suppose that the emergency occurs at some time t (according to density function $g(t) = \lambda e^{-\lambda t}$) between the j^{th} and $(j+1)^{st}$ inspection times (i.e., $t_j < t < t_{j+1}$) and the component is still operating at time t_j (which occurs with probability $1 - F(t_j)$). Then the penalty cost is incurred if the failure happens in time interval (t_j, t), which occurs with conditional probability of $\frac{F(t)-F(t_j)}{1-F(t_j)}$. Letting θ be the discount factor, it follows that the discounted probability of incurring the penalty cost is

$$\begin{aligned}\mathcal{A}_1(\mathcal{T}_n) &:= \sum_{j=0}^{n} (1 - F(t_j)) \int_{t_j}^{t_{j+1}} \frac{F(t) - F(t_j)}{1 - F(t_j)} \lambda e^{-\lambda t} e^{-\theta t} dt \\ &= \sum_{j=0}^{n} \int_{t_j}^{t_{j+1}} (F(t) - F(t_j))\, \lambda e^{-\lambda t} e^{-\theta t} dt, \end{aligned}$$

where $t_0 = 0$. The second expression, $(r + \mathcal{C}(\mathcal{T}_n))\mathcal{A}_2(\mathcal{T}_n)$, represents the discounted expected cost if either one of two replacement scenarios occur: a failure is detected by inspections or a replacement takes place at the scheduled time t_{n+1}. For $j \leq n$, we detect a failure at time t_j if the component has failed in interval (t_{j-1}, t_j) and the emergency has not occurred by the time of j^{th} inspection. Furthermore, a preventive replacement is performed at time t_{n+1}, if the component is still functional at the time of last inspection t_n and the emergency happens after the time of replacement t_{n+1}. Then, similar to the previous case, we can write

$$\mathcal{A}_2(\mathcal{T}_n) := \Big[\sum_{j=1}^{n} \big(F(t_j) - F(t_{j-1})\big) e^{-\lambda t_j} e^{-\theta t_j} \Big] + \big(1 - F(t_n)\big) e^{-\lambda t_{n+1}} e^{-\theta t_{n+1}}.$$

The third expression, $(r + \mathcal{C}(\mathcal{T}_n))\mathcal{A}_3(\mathcal{T}_n)$, corresponds to the discounted expected cost if the component has not failed by the time of emergency, in which case the emergency prompts the use of a good component, which needs restoration/replacement after use. Similar to the derivation of $\mathcal{A}_1(\mathcal{T}_n)$, one can show that

$$\mathcal{A}_3(\mathcal{T}_n) := \sum_{j=0}^{n} \int_{t_j}^{t_{j+1}} \big(1 - F(t)\big)\lambda e^{-\lambda t} e^{-\theta t} dt$$

Finally, the fourth expression, $c_i \mathcal{A}_4(\mathcal{T}_n)$, is the total discounted expected inspection cost in one cycle induced by policy $\mathcal{T}_n$. We derive this using the expected sum of indicator functions $\sum_{j=1}^{n} E[I_j]$, where each $I_j = 1$ if the j^{th} inspection is performed and 0 otherwise. Then, the expected number of inspections performed is equal to the sum of the probabilities that each inspection is performed for $j = 1, \ldots, n$. We perform the j^{th} inspection if the component has not failed by the time of $(j-1)^{th}$ inspection and the emergency has not occurred by the time of j^{th} inspection. The corresponding probability associated with this event is $(1 - F(t_{j-1}))e^{-\lambda t_j}$. Therefore, we have

$$\mathcal{A}_4(\mathcal{T}_n) := \sum_{j=1}^{n} \big(1 - F(t_{j-1})\big) e^{-\lambda t_j} e^{-\theta t_j}.$$

Now, we can rewrite (3.1) as follows:

$$\mathcal{C}(\mathcal{T}_n) = \frac{(c_p + r)\mathcal{A}_1(\mathcal{T}_n) + r\mathcal{A}_2(\mathcal{T}_n) + r\mathcal{A}_3(\mathcal{T}_n) + c_i \mathcal{A}_4(\mathcal{T}_n)}{1 - \mathcal{A}_1(\mathcal{T}_n) - \mathcal{A}_2(\mathcal{T}_n) - \mathcal{A}_3(\mathcal{T}_n)} = \frac{\mathcal{N}(\mathcal{T}_n)}{\mathcal{D}(\mathcal{T}_n)}, \quad (3.2)$$

where $\mathcal{N}(\mathcal{T}_n)$ and $\mathcal{D}(\mathcal{T}_n)$ are defined as the numerator and the denominator of the ratio after the first equality sign, respectively. Our goal is then to find an inspection policy that minimizes the above ratio. To this aim, we define:

$$\mathcal{L}_\alpha(\mathcal{T}_n) := \mathcal{N}(\mathcal{T}_n) - \alpha \mathcal{D}(\mathcal{T}_n), \quad \text{and} \quad (3.3)$$

$$\mathcal{R}(\alpha) := \min_{\mathcal{T}_n} \ \mathcal{L}_\alpha(\mathcal{T}_n). \quad (3.4)$$

Our problem is then equivalent to finding an α^* such that $\mathcal{R}(\alpha^*) = 0$, since it implies that for any $\mathcal{T}_n$, $\mathcal{N}(\mathcal{T}_n) - \alpha^* \mathcal{D}(\mathcal{T}_n) \geq \mathcal{R}(\alpha^*) = 0$ (i.e., $\mathcal{C}(\mathcal{T}_n) = \frac{\mathcal{N}(\mathcal{T}_n)}{\mathcal{D}(\mathcal{T}_n)} \geq \alpha^*$). It follows that $\mathcal{T}_n^*(\alpha^*)$, the minimizer of (3.4) for α^*, is the optimal inspection policy that minimizes (3.2) (because $\frac{\mathcal{N}(\mathcal{T}_n^*(\alpha^*))}{\mathcal{D}(\mathcal{T}_n^*(\alpha^*))} = \alpha^*$). There exists a unique α^* that solves $\mathcal{R}(\alpha) = 0$, since $\mathcal{R}(0) > 0$, and $\mathcal{R}(\alpha) < 0$ for any $\alpha > \frac{\mathcal{N}(\mathcal{T}_n)}{\mathcal{D}(\mathcal{T}_n)}$ (where $\mathcal{T}_n$ is any arbitrary inspection policy), and $\mathcal{R}(\alpha)$ is continuous and strictly decreasing in α (since $\frac{\partial \mathcal{R}(\alpha)}{\partial \alpha} = -\mathcal{D}(\mathcal{T}_n^*(\alpha)) < 0$ by the envelope theorem and since $\mathcal{D}(\mathcal{T}_n)$ is strictly positive due to discounting).

Therefore, we can find α^* by performing a simple binary search. Then, for each candidate value $\alpha = \tilde{\alpha}$, we need to find the optimal inspection policy $\mathcal{T}_n^*(\tilde{\alpha})$ that minimizes $\mathcal{L}_{\tilde{\alpha}}(\mathcal{T}_n)$. In Section 3.4, we discuss the properties of this minimization problem and provide an algorithm to obtain $\mathcal{T}_n^*(\tilde{\alpha})$.

We remark that the above approach for minimizing (3.2) is referred to as λ-minimization and was previously discussed and studied in Aven and Bergman [8].

3.4 Structural Results

In this section, we provide several properties of the optimal inspection policy and use these properties to develop an algorithm for obtaining the optimal inspection policy. We remark that since the optimal inspection policy for (3.2) is also the optimal inspection policy of (3.4) for α^*, it also satisfies all the properties presented in this section. Also, in this section, we treat α as a parameter and discuss solving the problem (3.4) for a fixed value of α. Therefore, we do not show the dependence of the optimal inspection policies on α in our notation. Before characterizing and identifying an optimal inspection policy, one first must establish its existence (later we discuss uniqueness as well). Proposition 3 in Appendix A shows the existence of an optimal inspection policy for problem (3.4) for any given value of α. We denote this policy by $\mathcal{T}_n^*$.

Now we consider the first order optimality condition. The function $\mathcal{L} = \mathcal{L}_\alpha$ defined in (3.3) is differentiable with respect to each t_j since $F(s)$ is a differentiable function (differentiability of $F(s)$ follows from the continuity of $f(s)$). The following lemma establishes that the optimal inspection policy must satisfy the first order necessary condition and cannot be a boundary solution (i.e., $0 < t_1 < t_2 < \ldots < t_{n-1} < t_n < t_{n+1}$). The first order necessary condition provides a tool for characterizing the optimal inspection policy. To be more specific, using the necessary condition, we are able to reduce the dimension of the decision variable space from $n+1$ to 1, which we discuss more following Lemma 1. It is first useful to define:

$$\Omega(t,x,y) := \Big[\frac{F(t)-F(t-x)}{f(t)} - \frac{1-e^{-(\lambda+\theta)y}}{\lambda+\theta}\Big] - \frac{c_i}{k}\Big[\frac{1-F(t)}{f(t)} + \frac{1}{\lambda+\theta}\Big], \tag{3.5}$$

where $k = \frac{\lambda}{\lambda+\theta}(c_p+r+\alpha)-(c_i+r+\alpha)$. We assume that $k > 0$ for all instances of the problem (3.4) that we solve. This assumption holds when the penalty cost c_p is significantly larger than the inspection cost c_i and the replacement cost r, or if the discount factor θ is small. In most practical settings, this assumption must hold to justify performing inspections and replacements. Function $\Omega(t, x, y)$ will be used in Lemma 1, as well as in proving several properties of the optimal inspection policy. Also, for ease of representation, we define the length of the j^{th} inspection interval as $x_j := t_j - t_{j-1}$.

Lemma 1 *Any optimal inspection policy $\mathcal{T}_n^* = \{t_j\}_{j=1}^{n+1}$ satisfies $0 < t_1 < \ldots < t_n < t_{n+1} < \infty$. Furthermore, necessary conditions for the optimality of an inspection policy are*

$$\Omega(t_j, x_j, x_{j+1}) = 0 \qquad \text{for each } 1 \le j < n, \tag{3.6}$$

$$\Omega(t_n, x_n, x_{n+1}) = -\frac{c_i e^{-(\lambda+\theta)x_{n+1}}}{k(\lambda+\theta)}, \tag{3.7}$$

$$\frac{F(t_n + x_{n+1}) - F(t_n)}{1 - F(t_n)} = \frac{\theta(r+\alpha)}{\lambda c_p}. \tag{3.8}$$

The necessary conditions (3.6-3.8) produce a relationship among the three subsequent values of t_j's, i.e., $\{t_{j-1}, t_j, t_{j+1}\}$. Since $t_0 = 0$, we can use (3.6) and (3.7) to find recursively all the inspection times by just determining the optimal first inspection time t_1. Therefore, for inspection policies that satisfy the necessary conditions (3.6-3.8), the inspection times are implicitly functions of $t_{1,n}$ (note we use the second subscript n to denote the total number of inspection opportunities available); thus, we denote the inspection times by $t_j[t_{1,n}]$ and the length of the inspection intervals by $x_j[t_{1,n}]$ to indicate this dependence.

Now we use properties of PF_2 densities (Lemma 9 of Appendix A) to develop some useful insights on the optimal inspection policy. Intuitively, the increasing failure rate property of PF_2 densities should induce decreasing lengths of inspection intervals as the component ages. In fact, other papers in the reliability literature have shown that when there are an unlimited number of inspection opportunities, the lengths of the inspection intervals are indeed decreasing in other contexts (e.g., Barlow et al. [11], Sengupta [68] and Parmigiani [54]). In contrast to these papers, our model considers a random variable for the time of emergency in addition to the failure time, and schedules only a finite number of inspections, after which, the component can be replaced. We will demonstrate in the numerical results of Section 3.5 that in our setting, inspection intervals may decrease at first but increase for later inspection intervals. The intuition is that in certain settings, one might use the inspection opportunities more conservatively when fewer of them remain and during periods where it is less likely to experience an emergency. This is formalized through the following theorem, which indicates that if there is an increase in duration from one interval to the next, the remaining intervals continue to increase.

Theorem 3 *Suppose that* $\mathcal{T}_n^* = \left\{t_j[t_{1,n}^*]\right\}_{j=1}^{n+1}$ *is an optimal inspection policy. If there exists* $m \in \{1, \ldots, n-1\}$ *such that* $x_{m+1}[t_{1,n}^*] > x_m[t_{1,n}^*]$*, then* $x_{j+1}[t_{1,n}^*] > x_j[t_{1,n}^*]$ *for all* $m \leq j \leq n$*.*

Note that Theorem 3 does not imply that the optimal inspection intervals necessarily increase; just that if there is an increase, then all subsequent intervals are also increasing in length.

Now we establish other properties of the inspection policies that satisfy the necessary conditions (3.6-3.8). We mentioned that for these inspection policies, the inspection times are implicitly functions of $t_{1,n}$. Our next lemma states how $t_j[t_{1,n}]$ and $x_j[t_{1,n}]$ change as functions of $t_{1,n}$. This sensitivity result will also be used to show that the optimal inspection policy is unique. Furthermore, this result plays a crucial role in developing the solution algorithm for finding the optimal inspection policy.

Lemma 2 *For any* $1 \leq j \leq n+1$*,* $x_j[t_{1,n}]$ *and* $t_j[t_{1,n}]$ *are strictly increasing in* $t_{1,n}$*.*

We can use Lemma 2 to prove that a unique value of $t_{1,n}$ satisfies the necessary conditions (3.6-3.8), and therefore, the optimal inspection policy is unique. Our solution algorithm is based on finding this unique optimal value of $t_{1,n}$.

Proposition 1 *The optimal inspection policy* $\mathcal{T}_n^* = \left\{t_j[t_{1,n}^*]\right\}_{j=1}^{n+1}$ *is unique for all* n*.*

Now we briefly discuss an algorithm for finding the optimal inspection policy (details of the algorithm can be found in Appendix E). Based on the discussion following Lemma 1, finding the optimum first inspection time $t_{1,n}^*$ suffices to determine the entire optimum inspection policy by using (3.6) and (3.7) (note that by Proposition 1, the optimal inspection policy is unique). Our algorithm uses Lemma 2 to perform a one-dimensional binary search to determine the optimal first inspection time $t_{1,n}^*$. The rest of the inspection times can then be determined using (3.6) and (3.7).

3.5 Numerical Examples

In the previous section, we established properties of the optimal inspection policy and outlined how we can use them to develop a solution algorithm. In this section, we explore the effect of changes in the model parameters as well as the effect of the availability of more inspection opportunities. Finally, we compare the performance of the optimal inspection policy with the best constant interval inspection policy (where the inspection times are scheduled equally-spaced from each other), which is easier to compute and implement.

To perform our experiments, we used Matlab 7.13 to implement our solution algorithms. It took negligible time to find the optimal inspection policy in all of our experiments. For our base-case scenario, we assume c_p, c_i and r are equal to 1000, 1 and 10, respectively. Furthermore, we assume that the emergency arrival time has an Exponential density with rate $\lambda = 0.2$ (i.e., $g(t) = 0.2e^{-0.2t}$), and the component fails according to a Weibull density function ($f(s) = \frac{\gamma}{\beta}\left(\frac{s}{\beta}\right)^{\gamma-1} e^{-(s/\beta)^{\gamma}}$), where $\gamma = 2$ and $\beta = 5$. These instances of the emergency arrival and component failure densities have (mean, standard deviation) of $(5, 5)$ and $(4.46, 1.62)$, respectively. Note that the Weibull distribution is commonly used to model failure times and for $\gamma = 1$, it coincides with the Exponential distribution. It is also known that the Weibull density is PF_2 for values of $\gamma \geq 1$.

3.5.1 Effect of changes in the number of available inspection opportunities

In this section, we explore the impact of having access to more inspection opportunities. In other words, we investigate how the optimal inspection policy $\mathcal{T}_n^*$ and the optimal discounted expected total cost $\mathcal{C}(\mathcal{T}_n^*)$ change as the limit n on the number of available inspection opportunities per cycle increases.

In Figure 3.1a, we see that the discounted expected total cost corresponding to the optimal inspection policy for the base-case strictly decreases at first when more inspection opportunities are available, but it becomes strictly increasing for $n \geq 9$. The cost becomes increasing because as the component ages, it requires more frequent inspections to protect against the catastrophe, which is costlier than replacing the component. Therefore, in this example, it is optimal to have nine inspection opportunities per cycle.

Figure 3.1b depicts the optimal inspection policy for different values of n, where each dashed curve corresponds to one value of n, and the star

Figure 3.1: The optimal inspection policy and the optimal discounted expected total cost (DETC) for different values of n ($\bigstar$ represents the replacement interval $t_{n+1} - t_n$).

(a) Optimal DETC for different n

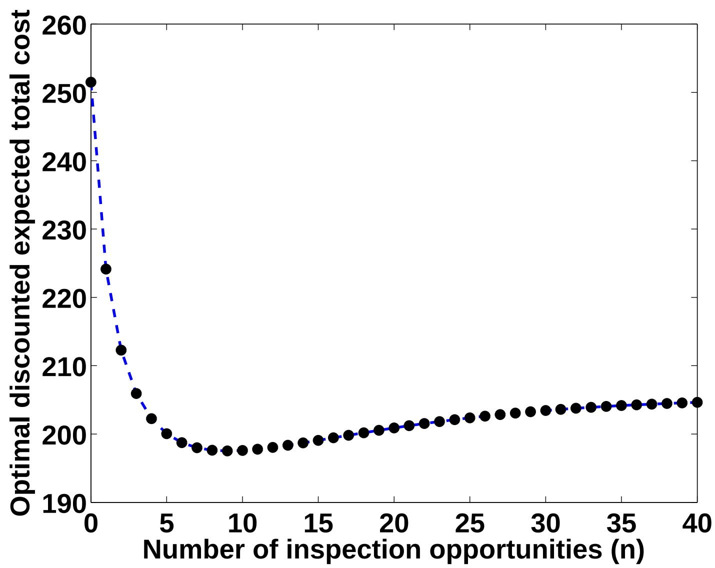

(b) Optimal inspection policies for different n

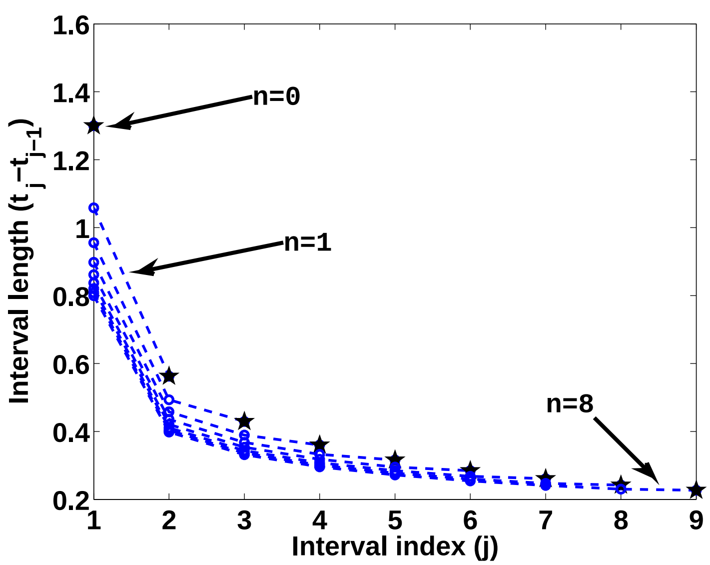

corresponds to the replacement interval after the last inspection. We observe in Figure 3.1b that for fixed n, the length of the inspection intervals are strictly decreasing (i.e., inspections are scheduled closer to each other as the component ages), which is consistent with the increasing failure rate property of the failure density. However, this need not hold for all sets of parameter values. If we change γ from 2 to 1.2, Figure 3.2a shows how for different values of n, inspection intervals are strictly decreasing at first, and whenever the lengths of the inspection intervals become strictly increasing, they remain strictly increasing, as predicted in Theorem 3. The reason for this behavior is that the failure rate increases at a slower rate when $\gamma = 1.2$ (compared to the base-case of $\gamma = 2$), and the optimal inspection policy is mostly derived by the emergency time distribution, meaning that inspections are scheduled during periods where it is more likely that an emergency occurs sooner rather than later when a failure is more likely.

We also observe in Figure 3.1b that for a fixed j, the length of the optimal inspection interval $x^*_{j,n}[t^*_{1,n}] = t^*_{j,n}[t^*_{1,n}] - t^*_{j-1,n}[t^*_{1,n}]$ is strictly decreasing in n, meaning that we schedule inspections closer to each other as more inspections become available. While this may also seem intuitive, Figure 3.2b demonstrates that this also is not a necessary condition of an optimal inspection policy. In this example, c_i is increased from 1 to 8 (all other baseline parameters being the same), which is close to the replacement cost ($r = 10$). Consequently, performing inspections and postponing the replacement is not as beneficial as before. Therefore, the optimal inspection policy schedules the inspection times farther apart to decrease the expected inspection cost because the resulting savings justify a slight increase in the probability of catastrophe.

3.5.2 Effect of changes in the failure time distribution

For our second experiment, we evaluate how sensitive our results are to changes in the component failure time density. In particular, we consider how changes in the probability density f interacts with g to affect cost. We do this by fixing the rate of emergency probability density λ at 0.2 and varying the parameter β from the Weibull distribution for the component failure time. In Figures 3.3a, 3.3b and 3.3c, we plot the discounted expected total cost, the probability of catastrophe in each cycle, and the expected number of inspections performed in each cycle, respectively, induced by the optimal inspection policy (for the case $n = 9$). However, instead of using β for the horizontal axis, we use a more intuitive indicator, p, which is defined as the probability that the failure occurs before the emergency

Figure 3.2: Examples of the optimal inspection policy with counter-intuitive properties.

(a) The lengths of the inspection intervals can be increasing for fixed n (for $\gamma = 1.2$).

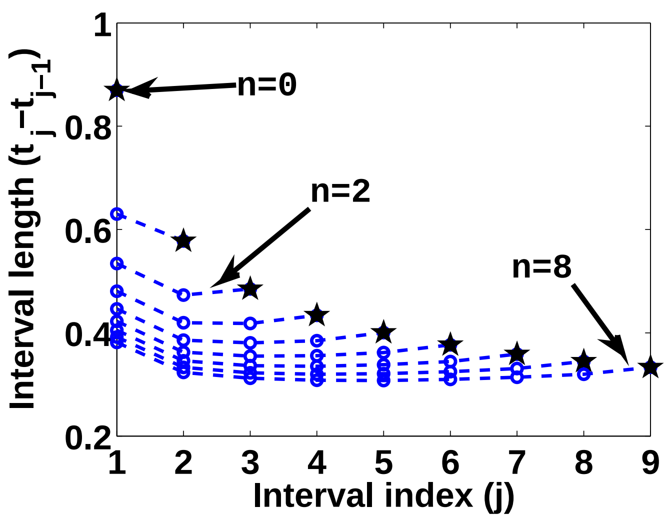

(b) The inspection times can be scheduled farther apart with more opportunities (for $c_i = 8$).

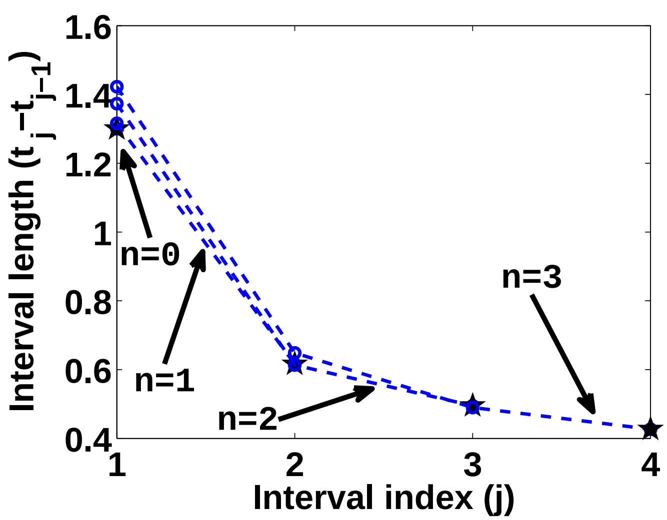

Figure 3.3: Three outcome measures when $p = prob(S < T)$ varies between 0 and 1 (■ corresponds to the base-case scenario).

(a) Optimal discounted expected total cost

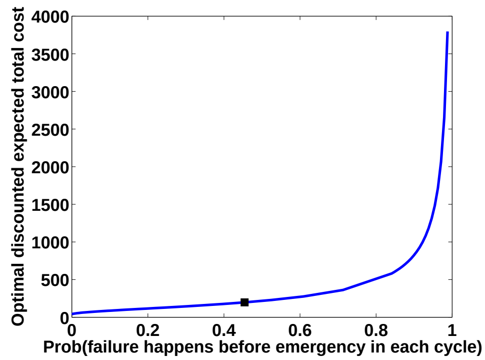

(b) Probability of catastrophe in each cycle

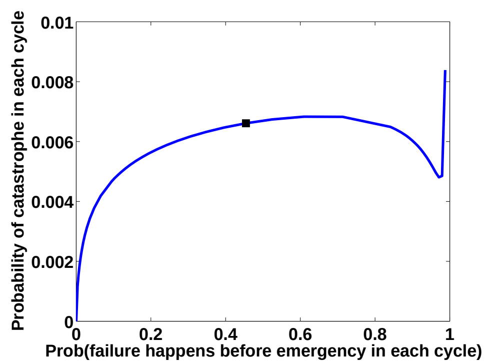

(c) Expected number of inspections performed in each cycle

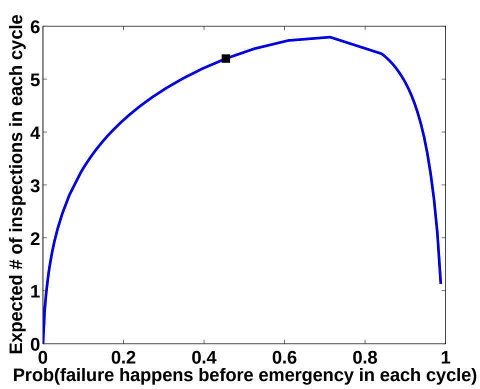

(i.e., $prob(S < T)$). For a constant value of the emergency arrival rate λ, this probability approaches 0 if $\beta \to \infty$, and it approaches 1 if $\beta \to 0$. In one extreme, when the component has a relatively long lifetime (i.e., as $p \to 0$), the discounted expected total cost of the optimal inspection policy $\mathcal{C}(\mathcal{T}_9^*)$ approaches 40. This cost only includes the discounted expected total replacement cost from the infinite stream of replacements after using a functioning component upon emergency. This is reasonable since even without performing any inspections, the probability that the component fails before the emergency is very small. Figures 3.3b and 3.3c also confirms this observation.

On the other extreme, if the failure occurs before the emergency with a probability close to 1 (i.e., the component has a short life), the discounted expected total cost is the largest. Again, this cost is also mainly derived by the cost of the replacements, as the probability of catastrophe is relatively small. In this case, the replacements are done after detecting the failure through the first inspection (note that the expected number of inspections performed per cycle approaches 1 as $p \to 1$). Since the component has a short lifetime (close to 0), one can wait and identify the failure by performing a single properly timed inspection. However, due to the component's short life, more frequent replacements are necessary in this case to provide coverage against the catastrophe. Consequently, as the component's lifetime decreases, the discounted expected total cost increases. We remark that for $0.98 < p < 1$, the component's extremely short lifetime makes it no longer economical to keep and maintain; therefore, we observe an increase in the probability of catastrophe (Figure 3.3b) because we are willing to accept a higher risk of incurring the penalty cost due to the high maintenance cost. Note that our base-case scenario ($\beta = 5$) results are shown by black squares.

In Figure 3.3c, we observe that for mid-range values of p, for which the chance of failure happening before emergency is close to the chance of emergency happening before failure, the expected number of inspections performed in each cycle is at its peak. For this range of p values, since it is harder to predict whether the emergency will happen before or after the failure, it is harder to guard against a catastrophe from occurring, despite optimally timing the inspection intervals (note that in Figure 3.3b, the probability of the catastrophe also has a peak for the intermediate values of p).

3.5.3 Effect of changes in the cost parameters

Finally, we explored the effect of changes in the cost parameters. To this aim, we consider the replacement cost r to be fixed, and vary c_i and c_p. Figure 3.4a depicts the optimal discounted expected total cost for $c_i = 1, 2, 5$ and 9. As expected, as the inspection cost increases and approach r, it is optimal to schedule smaller number of inspections before scheduling a replacement. In fact, for $c_i = 5$ and 9, our results suggest that we replace the component without scheduling any inspections because the inspection cost for these cases are close to the replacement cost. In contrast, when we vary the penalty cost c_p, our experiments does not reveal any specific pattern on the optimal number of inspection opportunities to schedule before the replacement (it is always optimal to schedule 9-11 inspections before the replacement). However, we observe that for the larger penalty cost, the magnitude of decrease in the discounted expected total cost is larger for the first few inspections opportunities.

Also, as expected, the discounted expected total cost increases as any of the cost parameters increase. Returning to our discussion in Section 3.5.2, the magnitude of this change depends on the value of p (i.e., probability of failure happening before emergency). As we mentioned earlier, in the two extremes, the optimal discounted expected total cost consists of mainly the expected total replacement cost. Therefore, the optimal discounted expected total cost is the most sensitive to the replacement cost r. On the other hand, for mid-range values of p, the penalty cost has a larger contribution to the total cost, and therefore its changes can also affect the total cost.

3.5.4 Comparison with constant interval inspection policies

In this experiment, we compare the performance (the discounted expected total cost) of the optimal inspection policy $\mathcal{T}_n^*$ with the best constant interval inspection policy. The constant interval inspection policy may be easier to compute and implement in practice. However, since this inspection policy is no longer optimal, it is important to know the performance loss associated with implementing it. First, we define an inspection policy with constant inspection interval of length T as $\widetilde{\mathcal{T}}_n := \left\{j \cdot T\right\}_{j=1}^{n+1}$, where n is the limit on the number of available inspection opportunities as before. In other words, we schedule n inspections at times $T, 2T, \ldots, nT$, and a replacement at time $(n+1)T$. For this experiment, we obtained $\widetilde{\mathcal{T}}_n^*$ by minimizing $\mathcal{C}(\widetilde{\mathcal{T}}_n)$ over the single variable T using Matlab's built-in solver.

Figure 3.5a shows the discounted expected total cost $\mathcal{C}(\widetilde{\mathcal{T}}_n)$ as a function

Figure 3.4: Optimal discounted expected total cost (DETC) for different n as c_i and c_p change.

(a) DETC for $c_i = 1, 2, 5$ and 9.

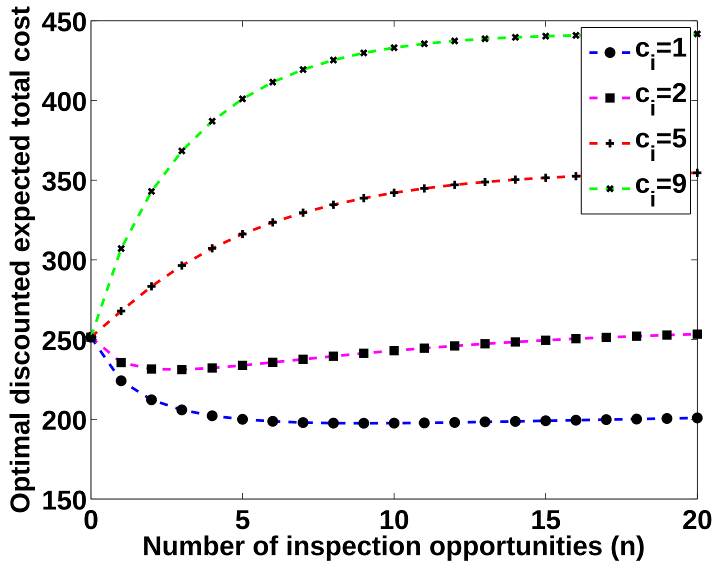

(b) DETC for $c_p = 50, 200, 600$ and 1000.

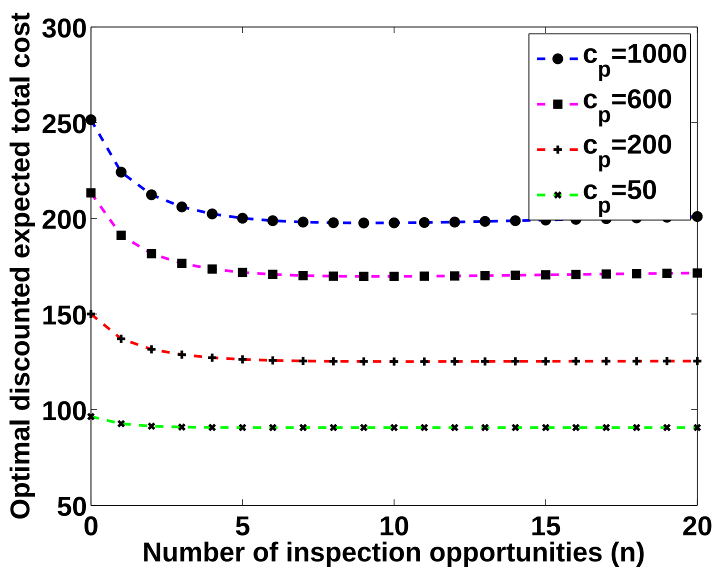

Figure 3.5: Comparison of constant interval inspection policies and optimal inspection policies.

(a) Discounted expected total cost as a function of T for different values of n

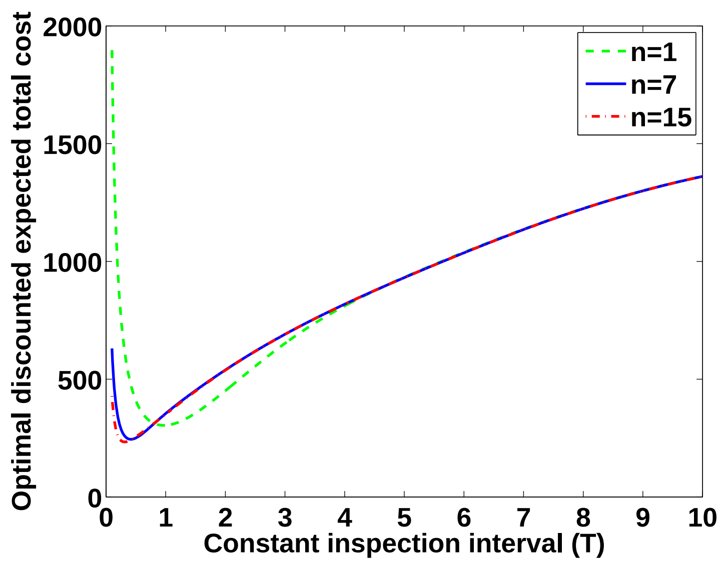

(b) Performance of optimal policies vs. best constant interval policies

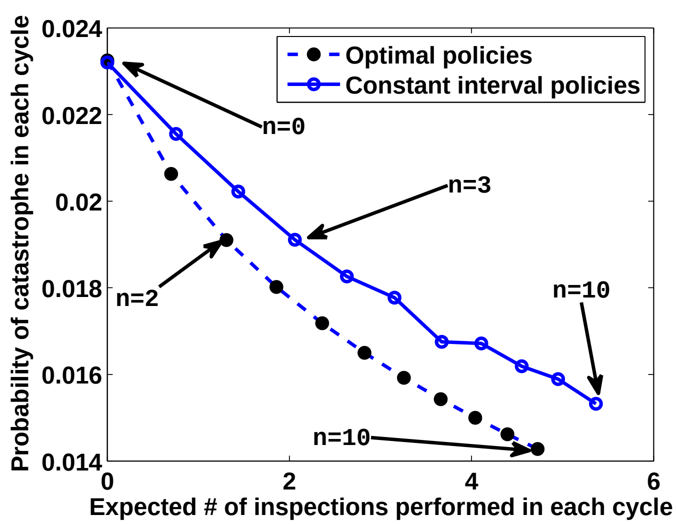

of constant interval length T for the values of $n = 1, 7$ and 15. For the baseline parameters and $n = 15$, the best constant interval policy is attained for $T^* = 0.36$ with the discounted expected total cost of 233.38, which is 18.2% more than the optimal expected total cost ($\mathcal{C}(\mathcal{T}_9^*) = 197.53$).

Figure 3.5b depicts two outcome measures: the probability of catastrophe and the expected number of inspections performed in each cycle both for the optimal inspection policies (dashed line) and the constant interval inspection policies (solid line) for the values of $n = 0, 1, \ldots, 10$. We observe that for $n = 0$, both policies perform exactly the same since in this case, the constant interval property does not impose any additional constraints. However, we observe as n increases, for the same probability of catastrophe, the expected number of inspections performed is much smaller if we space the inspections optimally. This means that the constant interval policies can maintain the same probability of catastrophe only at the expense of a larger expected number of inspections performed.

3.6 Conclusions

This chapter provides a new modeling framework for determining the optimal inspection times for a reliability problem. We defined a penalty cost for the situation where the component is not ready at the time of an emergency, and we also considered that there is a limit on the number of inspections opportunities, after which a replacement can be scheduled. Under the assumption that the component failure satisfies a mild technical condition (i.e., the PF_2 distribution), we showed that the optimal inspection policy is unique, and we provided an algorithm for finding it. Moreover, we established several interesting properties of the optimal solution. We also provided intuitive explanations for changes in the inspection policy due to changes in different parameters and supported the intuition with numerical examples. Our experiments also highlighted some counterintuitive results, which we discussed as well.

There are a number of managerial insights derived from our analysis that can be of interest to practitioners. First, one can use our model to decide when it is beneficial to pay extra to increase the inspection resources available by comparing the reduction in expected cost for each marginal inspection opportunity with the cost of obtaining it. Also, our numerical examples suggest that more inspections are expected to be performed in situations where it is harder to predict whether the emergency happens before or after the failure (Figure 3.3c), and when inspections are cheaper relative to the replacement cost (Figure 3.4a). Furthermore, we showed in Theorem 3 and observed in our numerical examples, in certain situations, the relative positioning of the emergency and failure distributions can result in an inspection policy where the inspection intervals could become increasing over time. This result sharply contrasts with the existing literature in which infinite inspection opportunities lead to inspection intervals that are decreasing. Therefore, despite the intuition that a component with the IFR property should be inspected more frequently as time passes, this logic may no longer hold when limited inspection opportunities are available because we want to schedule the limited remaining inspections over a longer period to delay the scheduled replacement (since short replacement cycles result in increased expected replacement cost).

Our approach of considering limited number of inspection opportunities and allowing for a preemptive replacement is not specific to this problem, and it can be applied to many similar problems previously studied in literature. Moreover, as a direction for future research, it would be interesting to study other versions of the problem, such as the case where inspections may be inaccurate or the inspections/replacements may take non-negligible time.

Chapter 4

Issuing Policies for Hospital Blood Inventory

4.1 Introduction

Red blood cells (RBCs), which deliver oxygen to body tissues, are the most common blood product used in transfusions. Donated RBC units are perishable and have a shelf life of 42 days in the United States and Canada. Hospitals regularly receive RBC units of different ages (measured from the time of donation) from regional blood centers and store them locally. Later, when demand for these units arise, they are issued (allocated) according to an issuing policy that considers factors such as the compatibility with recipient, the age of RBC unit issued, and the current available inventory.

The efficiency of an issuing policy can be measured based on several key performance metrics. The first metric is the outdate rate, which is the proportion of supplied RBC units that are discarded because they have reached the maximum allowable age of 42 days. In addition to the costs associated with collecting, distributing, holding, and discarding these units, reducing the outdates rate is especially important when the overall supply of blood is not sufficient to meet the demand. Another metric is the shortage rate, or the proportion of demand satisfied from another source since the hospital has no on-hand inventory. Again, dealing with shortages can be costly or difficult in urgent cases. Finally, while transfusing an RBC within its the allowable shelf life is considered safe by the current standards, several recent studies suggest that the transfusion value of the RBC units deteriorate over time (Wang et al. [76]), and that transfusing older blood can increase the risk of complications, particularly for critically ill patients (Koch et al. [33], Zallen et al. [79]). For instance, Koch et al. [33] concluded in their study that in patients undergoing cardiac surgery, transfusion of older RBCs was associated with a significantly increased risk of postoperative complications as well as reduced survival.

The current practice at hospitals is to issue RBCs in order from oldest to youngest inventory in order to minimize the number of outdates and shortages cases. In the Operations Research literature, this issuing policy, referred to as FIFO, has been shown to be the optimal issuing policies for several objective functions including outdate and shortage rates (Pierskalla and Roach [58]). The underlying assumption for these models is that the quality of a RBC unit remains constant throughout the 42 days lifespan. However, as we discussed above, recent findings in medicine suggest that the age of transfused blood can affect health outcomes, with older blood contributing to more complications. Given that the impact of RBC's age cannot be ignored, FIFO is no longer the optimal issuing policy. In fact, considering the objective of minimizing the average age of transfused blood would imply issuing RBCs in order from youngest to oldest (i.e., a LIFO policy).

In response to these clinical findings, there is a recent interest in the clinical community to design issuing policies that balance the trade-off between the shortage (and/or wastage) and the average age of transfused blood. For instance, Atkinson et al. [7] proposed the following class of policies based on a single age threshold: transfuse the oldest blood that is younger than the threshold, and if there is no blood younger than the threshold, transfuse the youngest blood that is older than the threshold. They used a simulation model based on data from the Stanford University Medical Center and showed that with a threshold of 14 days they can reduce the age of transfused blood without significantly affecting the amount of wastage or shortage. They demonstrated how changing the threshold affects the mean age of blood transfused and shortage percentage.

Atkinson et al. [7] suggested a simple class of policies, which aims to find the right balance between conflicting objectives of reducing the age of transfused blood and shortage. However, they neither characterized an optimal issuing policy, nor provided a framework for obtaining one. In this chapter, we analyze this problem as a perishable inventory management problem using an infinite-horizon dynamic programming model. We assume that the supply and demand processes are exogenous and random. Then, the only way to reduce shortages is to avoid outdates, meaning that an issuing policy that minimizes shortages also minimizes the outdates. Therefore, we define our objective function by considering only two costs: (1) a penalty cost associated with each RBC unit transfused, which is proportional to the age of the unit at the time of issuance, and represents the expected cost of any complications after the transfusion. (2) a penalty for each unit shortage which is the cost of satisfying the demand from another source. Then, our goal is to minimize the total expected discounted cost by deciding which units of blood on-hand and of what age to issue to satisfy the demand in each period. This dynamic programming formulation suffers from the curse of dimensionality due to large state and decision spaces. In order to deal with the intractable number of states, we use methods from the approximate dynamic programming (ADP) literature. In particular, we approximate the value function of the dynamic programming with an affine combination of a set of basis functions. We assign different weights to these functions by solving a linear program. Then, we use this approximate value function to find the approximate optimal issuing policy in each period by solving a tractable integer program.

The remainder of the chapter is organized as follows. In the next section, we review the relevant literature. In Section 4.3, we present our dynamic programming formulation of the problem. In Section 4.4, we discuss our solution method based on the linear programming approach to ADP, and we develop an efficient algorithm that finds an approximate optimal issuing policy. In Section 4.5, we present our numerical results, which are based on data from a large hospital in British Columbia. Finally, we discuss concluding remarks and future work in Section 4.6.

4.2 Literature Review

Our research in this chapter is related to the literature on blood inventory management, which is part of a larger stream of research on perishable inventory management. We restrict our focus to the models that are closely related to our problem (i.e., issuing policies) and refer the reader to Nahmias [43], Nahmias [44], and Karaesmen et al. [28] for comprehensive reviews of the early and recent developments in the area of perishable inventory systems (e.g., models dealing with optimal ordering policies).

The focus of the perishable inventory management literature (including the models that consider the management of the blood products) is more on the ordering decisions rather than the issuing policies, which is in most cases assumed to be FIFO. The limited studies that study the optimal issuing policies are mostly for the situations where either FIFO or LIFO is the optimal issuance policy. For instance, Pierskalla and Roach [58] showed that FIFO is in fact the optimal issuing policy for the following three objective functions: (1) maximizing the value of all demands satisfied, (2) minimizing the total number of units backlogged, and (3) minimizing the total number of outdates. More recently, Haijema et al. [25] used a Markov Decision Process model and simulation to develop near-optimal order-up-to replenishment policies for a number of issuing policies. Moreover, Deniz et al. [19] provided heuristics for finding the joint replenishment-issuing policies for a product with two periods of lifetime.

Another stream of research uses queueing theory to analyze perishable inventory systems. The general approach of these models is to analytically describe different performance metrics under simple issuing policies (in most cases, only the FIFO policy). As an example, Kaspi and Perry [29] obtained performance metric such as the average number of units in the system and the distribution of time between outdates for the FIFO issuing policy (also, see Kaspi and Perry [30], Perry [57], and Parlar et al. [53] for several extensions to this work). In a more recent work, Abouee-Mehrizi et al. [1] used this approach to find the distribution of the age of transfused blood and the shortage rates for the threshold policy introduced by Atkinson et al. [7]. Using queueing theory to analyze issuing policies has a major drawback, which is the fact that it can only consider very simple issuing rules and therefore one cannot use it to study general state-dependent issuing policies.

In general, dynamic inventory management models that deal with optimal ordering decisions for the perishable products are known to be extremely complicated because the dynamic programming formulation of these problems suffer from the curse of dimensionality. As a consequence, except for very simple cases (such as the case of two-period lifetime studied by Nahmias and Pierskalla [45]), the main focus of the literature is on developing efficient heuristic for finding near optimal ordering policies (see Nahmias [40], Fries [22], Nahmias [41], Nahmias [42], Cooper [16] as examples).

4.2.1 Contribution to related literature

Our work contributes to the literature on the allocation policies of perishable products. To our knowledge, our model is the first work that provides a dynamic (i.e., state-dependent) issuing policy whereas previous literature considers only static policies such as FIFO, LIFO, or age-based threshold policies as in Atkinson et al. [7] and Abouee-Mehrizi et al. [1]. While the FIFO issuing policy may be optimal for certain objective functions (see Pierskalla and Roach [58]), it does not capture the trade-offs between the conflicting objectives of age and shortages. Also, as we show in this chapter, the static issuing policies such as the threshold policy proposed by Atkinson et al. [7] is no longer optimal when we try to balance the shortage rate and the average age of blood transfused simultaneously. We develop an efficient ADP algorithm to find approximate optimal issuing policies and show that they can outperfom other classes of policies. Furthermore, our numerical experiments reveal several managerial insights regarding efficient blood issuing policies. In particular, the approximate optimal issuing policy found by our ADP algorithm motivates a simple, easily implementable policy, which still outperforms other policies.

4.3 Model and Formulation

We formulate this problem as an infinite horizon, discounted Markov Decision Process. We assign a penalty cost to each RBC unit transfused, which is proportional to the age of the unit at the time of issuance. For instance, this penalty cost can represent the expected cost of any complications after the transfusion. Furthermore, we consider another penalty for each unit shortage, which is the cost of satisfying the demand from another source. We consider a single blood type meaning that we do not consider transfusion of blood from other compatible blood types. We assume that the supply of blood units (and their age) and demand in each period are exogenous and uncertain, but their probability distributions are available. Then, our goal is to minimize the total expected cost by deciding which units of blood on-hand and of what age to issue to satisfy the demand in each period. First we define the notation:

- $i \in \{1, \ldots, 42\}$: age of the blood.
- $p(i) = ic$: penalty of satisfying a unit of demand with blood of age i, where $c > 0$ denotes the per-unit-period penalty cost.
- ℓ: penalty of satisfying one unit of demand from an outside source, when a shortage occurs. We assume $\ell > 42c$.
- Q_i, D: random variables for the number of age i blood arrivals and the number of demand units, respectively. We assume that these random variables are independent and identically distributed (iid) for different periods. We let q_i and d indicate the realized arrival of blood of age i and the realized demand.
- s_i: total supply of age i blood before satisfying the demand. This includes the leftover from the previous period as well as the new arrivals.
- u_i: total supply of age at most i before satisfying the demand (i.e., $u_i = \sum_{j=1}^{i} s_j$).

- x_i: number of units of age i used to satisfy the demand in the current period ($0 \leq x_i \leq s_i$ for each i).
- λ: the discount factor.

In each period the sequence of events is as follows. We observe this period demand d and the current supply vector $(s_1, s_2, \ldots, s_{42})$, and decide which units, $(x_1, x_2, \ldots, x_{42}$, to issue in order to satisfy the demand. Any part of the demand that is not satisfied by the units on-hand are satisfied from a secondary source through a rush order. Then, the remaining units of age 42 are discarded and the remaining blood units age one period. The new supply of the blood of different ages $(q_1, q_2, \ldots, q_{42})$ arrives and are added to the leftover units, resulting in a total supply inventory of $(q_1, q_2 + s_1 - x_1, \ldots, q_{42} + s_{41} - x_{41})$ at the beginning of the next period. Finally, the next period's demand d' is observed. Now, we define the different components of the MDP formulation.

Decision epochs: We assume that issuing decisions are made daily, after the demand for that day is realized.

State space: The state of the system consists of this period's supply vector and demand. In other words, we represent the state vector as $\mathcal{S} = (s_1, s_2, \ldots, s_{42}, d)$, where each s_i denotes the available blood units of age i and d denotes the current demand.

Action space: In each decision epoch, we let integer vector $\mathcal{X} = (x_1, x_2, \ldots, x_{42})$ represent the number of units, x_i, of each age i blood to issue. Note that if the units issued do not satisfy the whole demand, then the remaining demand is met from a secondary source through a rush order, but the system incurs a shortage cost per unit. Furthermore, for each state $\mathcal{S}$, the set of feasible actions must satisfy the following constraints.

$$\begin{aligned} & x_i \leq s_i && \text{for } i = 1, \ldots, 42 \\ & \sum_{i=1}^{42} x_i \leq d \\ & x_i \geq 0, \ \text{integer} && \text{for } i = 1, \ldots, 42. \end{aligned} \tag{4.1}$$

The first set of constraints specify that we cannot issue units more than the available supply of each age, whereas the second constraint ensures that we cannot issue units more than the current demand.

Transition probabilities: At the end of each period, the new supply of blood $\mathcal{Q} = (q_1, \ldots, q_{42}$ arrives and the next period's demand d' is realized. We assume that there is a finite number K of possible scenarios for the realized new supply and demand. The probability for each scenario is known, which we represent by $Pr(\mathcal{Q}, d')$, and it is independent of the current state or our action. Then, we can define the transition probabilities as follows:

$$p(\mathcal{S}'|\mathcal{S}, \mathcal{X}) = \\ \begin{cases} Pr(\mathcal{Q}, d') & \text{if } \mathcal{S}' = (q_1, q_2 + s_1 - x_1, \ldots, q_{42} + s_{41} - x_{41}, d'); \\ 0 & \text{otherwise.} \end{cases} \tag{4.2}$$

Basically, (4.2) suggest that given the starting state $\mathcal{S}$ and the action $\mathcal{X}$, with probability $Pr(\mathcal{Q}, d')$, we will start the next period with inventory vector $(q_1, q_2 + s_1 - x_1, \ldots, q_{42} + s_{41} - x_{41})$ and we must meet the demand d'. Note that the remaining units of age 42, $s_{42} - x_{42}$, at this period are discarded and the remaining units, $s_i - x_i$ for $i < 42$, have aged by one period.

Immediate cost: The immediate cost is associated with satisfying this period's demand and is a function of the current state and action. We define it as follows:

$$C(\mathcal{S}, \mathcal{X}) = \sum_{i=1}^{42} icx_i + \ell(d - \sum_{i=1}^{42} x_i), \tag{4.3}$$

where the first part, $\sum\limits_{i=1}^{42} icx_i$, denotes the total penalty cost associated with issuing blood from units on-hand and the second part, $\ell(d - \sum\limits_{i=1}^{42} x_i)$, corresponds to the cost of obtaining the shortage units from a secondary source.

Optimality equation: First, we define the value function for state $\mathcal{S}$ as the optimal expected discounted cost-to-go over an infinite horizon if we start the current period with state $\mathcal{S}$. We denote this cost by $V(\mathcal{S})$, which should satisfy the Bellman equations:

$$V(\mathcal{S}) = \min_{\mathcal{X}} \left\{ C(\mathcal{S}, \mathcal{X}) + \lambda \sum_{\mathcal{S}'} p(\mathcal{S}'|\mathcal{S}, \mathcal{X}) V(\mathcal{S}') \right\} \qquad \forall \mathcal{S}. \tag{4.4}$$

Because of the large state and action spaces, we cannot use the traditional methods for solving MDPs, such as the value iteration or the policy iteration algorithms. For instance, if the demand and the number of units of each age never exceeds 100, the state space is of size 100^{43}.

4.4 Solution Approach

Since it is intractable to exactly solve the MDP formulation of Section 4.3, we solve the equivalent linear programming formulation through ADP. We first define the following approximation to the value function:

$$V(\mathcal{S}) \cong \widetilde{V}(\mathcal{S}) = \theta_0 + \sum_{i=1}^{42} \theta_i u_i + \delta d + \sigma[d - u_{42}]^+, \tag{4.5}$$

where $u_i := \sum_{j=1}^{i} s_j$ represents the total supply available of age at most i, and $[d - u_{42}]^+ = \max(0, d - u_{42})$ represents the shortage given the total available supply of u_{42} and the demand d. Then, we can use ADP algorithms to find the best set of coefficients $(\theta_0, ..., \theta_{42}, \delta, \sigma)$ such that (4.5) would be a good approximation for the exact value function. Note that each θ_i (if negative) can be viewed as the marginal savings in cost for each additional unit of age at most i, whereas δ can be interpreted as the marginal cost for each additional demand unit and σ as the marginal cost for each additional unit of shortage. By defining the approximation as above, we try to capture the main factors deriving the cost, i.e., the supply of different ages, the demand, and the magnitude of a shortage if it occurs.

In the remainder of Section 4.4, we discuss our approach for determining the best values for coefficients $(\theta_0, ..., \theta_{42}, \delta, \sigma)$ and how to use the resulting approximation to determine the issuing policy in each period. To this aim, we introduce several optimization problems, which we overview here. To tune the coefficients $(\theta_0, ..., \theta_{42}, \delta, \sigma)$, we solve the linear programming model described in (4.6) or its dual (4.7) using column generation. In order to generate new columns, we solve mixed integer programs of form (4.9). Finally, we solve problem (4.15) or its equivalent form (4.17) to determine the approximate issuing policy in each period.

4.4.1 Calibrating approximate value function coefficients

To tune the coefficients $(\theta_0, ..., \theta_{42}, \delta, \sigma)$, we use the linear programming approach to ADP (originally proposed by Schweitzer and Seidmann [67]). We first introduce the linear program form of our original MDP by rewriting the optimality equations (4.4) as follows.

$$\max \qquad \sum_{\mathcal{S}} \alpha(\mathcal{S}) V(\mathcal{S})$$

subject to

$$V(\mathcal{S}) \leq C(\mathcal{S}, \mathcal{X}) + \lambda \sum_{\mathcal{S}'} p(\mathcal{S}'|\mathcal{S}, \mathcal{X}) V(\mathcal{S}') \qquad \forall \mathcal{S}, \mathcal{X}.$$

In the above linear program, we have a decision variable $V(\mathcal{S})$ for each possible state vector $\mathcal{S}$. The optimal solution to the above linear program is the same as the solution to the optimality equations of the MDP provided in (4.4) given that the objective coefficients $\alpha(\mathcal{S})$ in the above linear program are all positive. It is convenient to normalize these coefficients such that $\sum_{\mathcal{S}} \alpha(\mathcal{S}) = 1$ and view them as a probability distribution over the initial state of the system (Puterman [61, p. 223]).

Note that due to the large state and action spaces, this linear program also suffers from curse of dimensionality, meaning it cannot be solved efficiently. To overcome the curse of dimensionality, we can replace $V(\mathcal{S})$ by our approximation value function defined in (4.5). The resulting problem is a linear programming problem that has 45 decision variables $(\theta_0, ..., \theta_{42}, \delta, \sigma)$:

$$\max \qquad \theta_0 + \sum_{i=1}^{42} \mathbb{E}_\alpha \left[u_i\right] \theta_i + \mathbb{E}_\alpha \left[d\right] \delta + \mathbb{E}_\alpha \Big[[d - u_{42}]^+ \Big] \sigma \tag{4.6}$$

subject to

$$(1-\lambda)\, \theta_0 + \sum_{i=1}^{42} \Theta_i(\mathcal{S}, \mathcal{X}) \theta_i + \Delta(\mathcal{S}) \delta + \varSigma(\mathcal{S}, \mathcal{X}) \sigma \leq C(\mathcal{S}, \mathcal{X}) \qquad \forall \mathcal{S}, \mathcal{X}$$

where

$$\mathbb{E}_\alpha \left[u_i\right] := \sum_{\mathcal{S}} \alpha(\mathcal{S}) u_i(\mathcal{S}) \qquad \forall i$$

$$\mathbb{E}_\alpha \left[d\right] := \sum_{\mathcal{S}} \alpha(\mathcal{S}) d(\mathcal{S})$$

$$\mathbb{E}_\alpha \Big[[d - u_{42}]^+ \Big] := \sum_{\mathcal{S}} \alpha(\mathcal{S}) [d(\mathcal{S}) - u_{42}(\mathcal{S})]^+$$

and

$$\begin{aligned}
\Theta_i(\mathcal{S},\mathcal{X}) &:= u_i(\mathcal{S}) - \lambda \sum_{(\mathcal{Q},d')} Pr(\mathcal{Q},d')u'_i(\mathcal{S},\mathcal{X},\mathcal{Q}) && \forall i \\
\Delta(\mathcal{S}) &:= d(\mathcal{S}) - \lambda \sum_{(\mathcal{Q},d')} Pr(\mathcal{Q},d')d' \\
\Sigma(\mathcal{S},\mathcal{X}) &:= [d(\mathcal{S}) - u_{42}(\mathcal{S})]^+ - \lambda \sum_{(\mathcal{Q},d')} Pr(\mathcal{Q},d')[d' - u'_{42}(\mathcal{S},\mathcal{X},\mathcal{Q})]^+.
\end{aligned}$$

Note that in (4.6), the expectations over vector α are estimated by sampling states (see Bertsekas and Tsitsiklis [13, p. 377] for more details). The above formulation reduces the number of decision variables significantly, but this problem still has a large number of constraints because we have one constraint for each possible state-action pair. Fortunately, for finding the optimal solution, it suffices to find at most 45 of these constraints which are binding at optimality (by the fundamental theorem of linear programming). Then, we can start with an initial set of 45 constraints and add new cuts (i.e., constraints) by finding the most violated constraint. We stop adding new cuts and report the solution when the optimality gap is smaller than a required precision. Equivalently, we can solve the dual of the above problem using delayed column generation. The dual to problem (4.6) can be written as follows:

$$\min \quad \sum_{\mathcal{S},\mathcal{X}} C(\mathcal{S},\mathcal{X})\mathcal{W}(\mathcal{S},\mathcal{X}) \tag{4.7}$$

subject to

$$\begin{aligned}
(1-\lambda)\sum_{\mathcal{S},\mathcal{X}} \mathcal{W}(\mathcal{S},\mathcal{X}) &= 1 \\
\sum_{\mathcal{S},\mathcal{X}} \Theta_i(\mathcal{S},\mathcal{X})\mathcal{W}(\mathcal{S},\mathcal{X}) &= \mathbb{E}_\alpha[u_i] && \forall i \\
\sum_{\mathcal{S},\mathcal{X}} \Delta(\mathcal{S})\mathcal{W}(\mathcal{S},\mathcal{X}) &= \mathbb{E}_\alpha[d] \\
\sum_{\mathcal{S},\mathcal{X}} \Sigma(\mathcal{S},\mathcal{X})\mathcal{W}(\mathcal{S},\mathcal{X}) &= \mathbb{E}_\alpha\Big[[d-u_{42}]^+\Big] \\
\mathcal{W}(\mathcal{S},\mathcal{X}) &\geq 0 && \forall \mathcal{S},\mathcal{X}.
\end{aligned}$$

The above problem has a dual variable for each state-action pair, but at most 45 of these variables are needed to be non-negative at optimality.

Therefore, we can use delayed column generation approach and add new variables as needed. Finding an initial set of columns (variables), which produce a feasible solution to problem (4.7) is not straightforward. Thus, we use the Phase I method of linear programming to find this initial set. The Phase I method starts with columns of an identity matrix of size 45, meaning that we add a non-negative slack variable to each constraint of the dual problem. Then, we try to minimize the sum of these slack variables by generating and adding columns from the original dual problem. If the Phase I objective function becomes zero, then we can remove the columns of the identity matrix and the remaining columns provide a feasible solution to the dual problem. We do not explain how we choose the columns that are added to the Phase I problem because it is similar to our approach for adding constraints to the primal problem, which we explain next.

Finally, having this initial set of columns for problem (4.7), we solve the relaxed dual problem (which we call the master problem) and find its optimal solution as well as the corresponding optimal solution to the relaxed primal problem (4.6) (i.e., $(\theta_0^*, ..., \theta_{42}^*, \delta^*, \sigma^*)$). Since we did not consider all the constraints from the original problem, this solution may not satisfy all the constraints to the primal problem. We find the most violated constraint by solving a pricing sub-problem that finds a feasible state-action pair $(\mathcal{S}, \mathcal{X})$ that maximizes

$$\widetilde{V}(\mathcal{S}) - \lambda \sum_{\mathcal{S}'} p(\mathcal{S}'|\mathcal{S}, \mathcal{X}) \widetilde{V}(\mathcal{S}') - C(\mathcal{S}, \mathcal{X}). \tag{4.8}$$

Recall that the primal constraints for problem (4.6) are of the form

$$\widetilde{V}(\mathcal{S}) \leq C(\mathcal{S}, \mathcal{X}) + \lambda \sum_{\mathcal{S}'} p(\mathcal{S}'|\mathcal{S}, \mathcal{X}) \widetilde{V}(\mathcal{S}') \qquad \forall \mathcal{S}, \mathcal{X}.$$

If the maximum of (4.8) is greater than 0 for the optimal state-action pair $(\mathcal{S}^*, \mathcal{X}^*)$, it means that the constraint corresponding to this state-action pair is the most violated constraint. On other hand, if the maximum is not positive, then the current solution $(\theta_0^*, ..., \theta_{42}^*, \delta^*, \sigma^*)$ satisfies all the constraints from the original primal problem.

Given $(\theta_0^*, ..., \theta_{42}^*, \delta^*, \sigma^*)$, we can represent the pricing sub-problem as the follows:

$$\max \quad (1 - \lambda)\, \theta_0^* + \sum_{i=1}^{42} \theta_i^* \Theta_i(\mathcal{S}, \mathcal{X}) + \delta^* \Delta(\mathcal{S}) + \sigma^* \mathit{\Sigma}(\mathcal{S}, \mathcal{X}) - C(\mathcal{S}, \mathcal{X}) \tag{4.9}$$

where

$$\begin{aligned}
\Theta_i(\mathcal{S},\mathcal{X}) &:= \sum_{j=1}^{i} s_j - \lambda \sum_{(\mathcal{Q},d')} Pr(\mathcal{Q},d') \left(\sum_{j=1}^{i} q_j + \sum_{j=1}^{i-1}(s_j - x_j) \right) && \forall i \\
\Delta(\mathcal{S}) &:= d - \lambda \sum_{(\mathcal{Q},d')} Pr(\mathcal{Q},d')d' \\
\varSigma(\mathcal{S},\mathcal{X}) &:= \left[d - \sum_{j=1}^{42} s_j \right]^+ - \lambda \sum_{(\mathcal{Q},d')} Pr(\mathcal{Q},d')[d' - u'_{42}]^+ \\
u'_{42} &:= \sum_{j=1}^{42} q_j + \sum_{j=1}^{41}(s_j - x_j).
\end{aligned}$$

In the sub-problem (4.9), the decision variables are the state $(s_1,\ldots,s_{42},d)$ and the action $(x_1,\ldots,x_{42})$. Therefore, to ensure their feasibility (meaning that the action vector does not issue blood units more than what is available according to the state vector), we consider only state-action pairs that satisfy the following set of linear constraints:

$$\begin{aligned}
x_i &\le s_i && \forall i \\
\sum_{i=1}^{42} x_i &\le d \\
x_i, s_i &\ge 0, \text{ integer} && \forall i \\
d &\ge 0, \text{ integer}\ .
\end{aligned}$$

Note that $\Theta_i(\mathcal{S},\mathcal{X})$, $\Delta(\mathcal{S})$, and all the constraints are linear in the decision variables. However, $\varSigma(\mathcal{S},\mathcal{X})$ is not a linear function of the decision variables because it contains piecewise linear functions of the decision variables. We convert this function into a linear function by defining integer variables k_0^+, k_0^-, $k^+(\mathcal{Q},d')$ and $k^-(\mathcal{Q},d')$, and binary variables b_0 and $b(\mathcal{Q},d')$:

$$\varSigma(\mathcal{S},\mathcal{X}) = k_0^+ - \lambda \sum_{(\mathcal{Q},d')} Pr(\mathcal{Q},d')k^+(\mathcal{Q},d'),$$

where k_0^+ satisfies the following set of linear constraints (letting N be a large positive integer):

$$k_0^+ + k_0^- - N = d - \sum_{j=1}^{42} s_j \tag{4.10}$$
$$k_0^+ \leq N b_0 \tag{4.11}$$
$$k_0^- \geq N b_0 \tag{4.12}$$
$$k_0^- \leq N \tag{4.13}$$
$$k_0^+, k_0^- \geq 0, \text{ integer and } b_0 \in \{0, 1\}.$$

Note that if $d - \sum_{j=1}^{42} s_j > 0$ (i.e., we face shortage), then the first constraint (4.10) together with (4.13) imply that $k_0^+ > 0$. Then, it follows from (4.11) that $b_0 = 1$. Then, we have $k_0^- = N$ from (4.12) and (4.13). Consequently, k_0^+ denotes the shortage because $k_0^+ = d - \sum_{j=1}^{42} s_j$ by (4.10). On the other hand, if $d - \sum_{j=1}^{42} s_j \leq 0$, then the above set of constraints ensure that $k_0^+ = 0$. Similarly, for each $(\mathcal{Q}, d')$, $k^+(\mathcal{Q}, d')$ must satisfy the following:

$$k^+(\mathcal{Q}, d') + k^-(\mathcal{Q}, d') - N = d' - \sum_{j=1}^{42} q_j + \sum_{j=1}^{41} (s_j - x_j) \tag{4.14}$$
$$k^+(\mathcal{Q}, d') \leq N b(\mathcal{Q}, d')$$
$$k^-(\mathcal{Q}, d') \geq N b(\mathcal{Q}, d')$$
$$k^-(\mathcal{Q}, d') \leq N$$
$$k^+(\mathcal{Q}, d'), k^-(\mathcal{Q}, d') \geq 0, \text{ integer and } b(\mathcal{Q}, d') \in \{0, 1\}.$$

Consequently, sub-problem (4.9) becomes an integer program, which in our case can be solved quickly. After solving the above integer program, the optimal values of the decision variables $(s_1^*, \ldots, s_{42}^*, d^*)$ and $(x_1^*, \ldots, x_{42}^*)$ provides us with the state-action pair that corresponds to the most violated constraint for the current solution to the master problem. We add this constraint to the master problem and resolve it to find a new solution. Then, we use the new solution to find another violated constraint. We stop adding more cuts when the termination criteria is met (i.e., when no more constraints are violated or when the optimality gap for the master problem is within the required range).

4.4.2 ADP-based issuance policy

Now, we explain how we can use the above approximation to determine an approximate optimal issuing policy in each period. Since it is impractical to store the approximate optimal issuing policy for each possible state, we use our approximate value function to find the approximate optimal issuing policy as needed for each state realized in practice.

In each period, given the state vector $\mathcal{S}$, we solve the following problem to find an approximate optimal issuing policy:

$$\min_{\mathcal{X}} \Big\{ C(\mathcal{S}, \mathcal{X}) + \lambda \sum_{\mathcal{S}'} p(\mathcal{S}'|\mathcal{S}, \mathcal{X}) \widetilde{V}(\mathcal{S}') \Big\}, \tag{4.15}$$

where given the initial state $\mathcal{S}$, our action $\mathcal{X}$, the new arrival vector $\mathcal{Q}$, and the new demand d', we have

$$\widetilde{V}(\mathcal{S}') = \theta_0 + \sum_{i=1}^{42} \theta_i u'_i + \delta d' + \sigma \left[d' - u'_{42} \right]^+ . \tag{4.16}$$

Note that in (4.16), each $u'_i = \sum_{j=1}^{i} q_j + \sum_{j=1}^{i-1} (s_j - x_j)$ is the total inventory of age at most i in the next period given the initial state $\mathcal{S}$, our action $\mathcal{X}$, and the new arrival vector $\mathcal{Q}$. This quantity is equal to the sum of new blood units arrived of age at most i, and the remaining aged units of age at most i after satisfying this period's demand.

A feasible issuing policy must satisfy the constraints in (4.1). Then, we can write the minimization problem (4.15) as follows using (4.3) to replace the immediate cost $C(\mathcal{S}, \mathcal{X})$:

$$\min \quad \sum_{i=1}^{42} icx_i + \ell(d - \sum_{i=1}^{42} x_i) + \lambda \sum_{(\mathcal{Q}, d')} Pr(\mathcal{Q}, d') \widetilde{V}(\mathcal{S}') \tag{4.17}$$

subject to

$$\begin{aligned} x_i &\leq s_i && \text{for } i = 1, \ldots, 42 \\ \sum_{i=1}^{42} x_i &\leq d && \\ x_i &\geq 0, \ \text{integer} && \text{for } i = 1, \ldots, 42 \end{aligned}$$

Note that the constraints of the above problem are all linear in the decision variables x_i. However, the objective function has a nonlinear part (i.e.,

$[d' - u'_{42}]^+$). We can use an approach similar to (4.14) to convert the above problem to an integer program. We solve this integer program in each period to find the approximate optimal issuing policy. In Section 4.5, we will evaluate the performance of the ADP-based policy and provide several insights based on our numerical experiments. Furthermore, we compare the performance of our policy with other classes of policies proposed in the literature.

4.5 Numerical Experiments

In this section, we use our solution approach discussed in Section 4.4 to find approximate optimal issuing policies for the case of a large hospital in British Columbia. We use historical data for this hospital to calibrate different parameters of our model. Then, using a simulation model, we evaluate the performance of our policy and provide several managerial insights on the properties of good issuing policies that balance the trade-offs between the age of blood transfused and the shortage rate. Furthermore, we show how our policy compares to issuing policies currently used in practice as well as previously proposed policies in the literature. Finally, using insights from our ADP-based issuing policy, we propose a simple static issuing policy that performs nearly as good as the ADP-based policy, but is easier to be implemented in practice. The simulation model was implemented in Matlab R2013a while the column generation algorithm and the integer programming models were implemented in AMPL with CPLEX 12.2 as the solver.

4.5.1 Simulation of the hospital blood bank

In order to evaluate the performance of different blood issuing policies, we developed a simulation model that mimics the dynamics of a hospital blood bank (the blood bank is a division of a hospital where the received RBC units are stored before being issued for transfusion). Our simulation model consists of four steps which we discuss next.

1. Demand realization: In each period, we randomly generate the daily demand according to an empirical distribution obtained from historical data for our hospital. In particular, we used the data on daily demand for a period of one year from April 1, 2010 to March 31, 2011. The demand and supply information for different blood types is not specified in our data. Therefore, we consider only a single blood type for this study.

2. Supply realization: We also generate the number of blood units arriving at the hospital each day, by age, using the empirical distribution based on historical data. In practice, the hospital receives blood from two different sources. The main source of blood supply is by ordering directly from Canadian Blood Services (CBS). In addition to receiving blood from CBS, the hospital receives blood units from smaller hospitals. These smaller hospitals send their blood units to the larger hospital when these units are within 10 days from expiration. The reason for this inter-hospital redistribution program is to reduce the wastage of blood since these units are more likely to be transfused at the larger hospital which faces a higher demand rate.

3. Issuing policy: The new blood arrived is added to our current inventory. Then, the simulation model passes the demand and total supply vector to the issuing policy module. For a given issuing policy, this module determines which units should be allocated to satisfy the demand and if the system faces a shortage. For instance, if the current issuing policy is simply FIFO, then the issuing policy allocates the oldest units on hand to satisfy the demand. On the other hand, if we want to test our ADP-based issuing policy, the simulation program calls our optimization model to determine which blood units should be issued to satisfy the demand.

4. Updating metrics: Finally, knowing the units issued on each day, we update our performance metrics. The simulation program gathers information regarding the shortage, wastage, and age of blood transfused as our main performance metrics. In calculating the average age of blood transfused, we do not consider the age of blood units received in case of a shortage. In this step, we also update the inventory of blood at the start of the following period. In other words, the blood units of age 42 are discarded and the remaining units age by 1 day.

4.5.2 Results and insights

In this section, we use our simulation program to evaluate the performance of the ADP-based issuing policy and compare it against other classes of policies. Moreover, we discuss several important insights regarding good issuing policies gleaned from our numerical experiments. Finally, we explore the sensitivity of our results with respect to changes in the blood supply and demand rates.

The performance metrics reported for each issuing policy are based on 100 simulation repetitions, where each repetition consists of simulating 1000 days of a hospital blood bank. We consider the first 700 days as the warm-up period and calculate all metrics using the information for the last 300 days.

For our baseline scenario, we set the per-unit-period penalty cost $c = 10$ and the shortage cost $\ell = 1000$. The values for the cost parameters are chosen arbitrarily and do not represent actual medical costs. In practice, the relationship between the age of blood transfused and the corresponding cost of complications is not well understood. In the absence of such data, we decided to find the approximate optimal issuing policies for different values of the cost parameters and depict the performance of these policies in the form of a trade-off curve of two conflicting objectives (the average age of blood transfused and the shortage rate). Then, policy makers can use this curve and choose the issuing policy that they believe balances these objectives in the best way.

Comparison of ADP-based and age-based threshold policies

The results of our first set of experiments are depicted in Figure 4.1. In this figure, each point represents a blood issuing policy. The dashed line corresponds to the threshold policies proposed by Atkinson et al. [7] and the solid line corresponds to our ADP-based issuing policies. The vertical axis denotes the average age of blood transfused, whereas the horizontal axis denotes the other performance metric of interest, the shortage rate.

The threshold policy of Atkinson et al. [7] issues blood unit based on the following rule: on each day, issue the oldest blood that is the same age or younger than the threshold, and if there is no blood younger than the threshold then issue the youngest blood that is older than the threshold. The different points on the dashed-line are associated with different choices of the threshold, where the upper-left point corresponds to the threshold of 42 days (which is equivalent to FIFO policy) and the lower-right point corresponds to the threshold of 1 day (i.e., LIFO issuing policy).

For the ADP-based policies on the solid line, different issuing policies are obtained by varying the cost ratio ℓ/c (in fact, we fix $c = 10$ and change ℓ). The lower-right point corresponds to the ADP-based issuing policy obtained for the case of $\ell = 420$, meaning that the shortage cost is the same as the penalty cost for issuing a blood unit of age 42. In this case, the issuing policy performs very similar to the LIFO policy because there is no incentive to avoid a shortage scenario by issuing the units which are going to expire soon. In contrast, as we increase the shortage cost to the other extreme ($\ell = 10,000$, upper-left point of the solid curve), the corresponding issuing policy resembles a FIFO issuing policy because it tries to avoid any shortage scenarios which are very costly.

When the shortage cost has a moderate value, the ADP-based issuing

Figure 4.1: Comparison of ADP-based policies and age-based threshold policies.

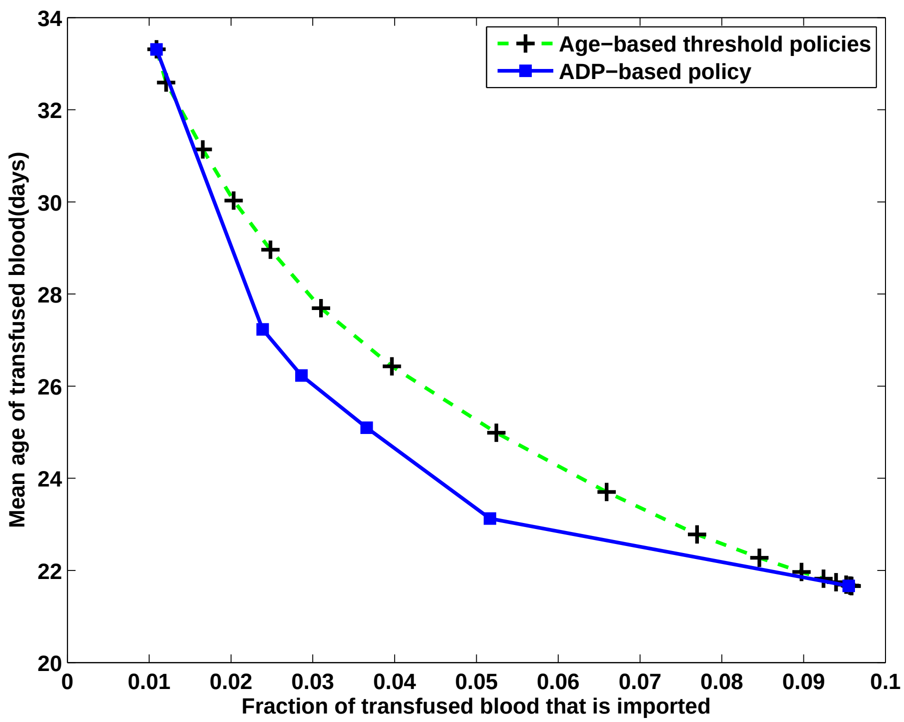

policy outperforms the threshold policies. For instance, for our baseline scenario ($\ell = 1000$), the average age of blood transfused for the ADP-based issuing policy is 25.10 and the shortage proportion is 0.037. If we issue the blood according to a threshold policy that has the same value of average age of blood transfused, then we face a 30% increase in the shortage proportion. Furthermore, for the shortage proportion of 0.052, the ADP-based policy results in 8% lower average age of blood transfused (23.13 vs. 25.10).

Our ADP-based policy is state-dependent, meaning that it considers the current state of the system in deciding which units to issue in each period. For instance, if the current on-hand inventory suggests that a shortage might occur soon, then the ADP-based policy would suggest issuing blood according to the FIFO policy, whereas it might suggest issuing fresher blood if we have excess supply. In contrast, the threshold policy is static and always use the same threshold regardless of the current state of the system. Atkinson et al. [7] showed that using their threshold policy with a carefully chosen threshold, one can reasonably decrease the average age of blood transfused with only a moderate increase in the shortage rate. Our results suggest that we can achieve the same average age of blood transfused with an even lower increase in the shortage rate if the issuing policy adapts to current state of the system.

Our results provide stronger evidence for policy makers that they can decrease the average age of blood transfused without jeopardizing the availability of the blood supply. Moreover, we showed that the state-dependent issuing policies perform better than the static policies such as FIFO, LIFO or the threshold policy. Our ADP-based policy smartly issues fresher blood when the likelihood of facing a shortage is low, and avoids outdates by issuing older blood units if the current on-hand inventory suggests possibility of a shortage scenario.

Sensitivity to supply-demand ratio

Our second set of numerical experiments explore the sensitivity of our results under different supply-demand regimes. To this aim, in addition to our base-line scenario where the annual total supply matches the annual total demand (i.e., supply-demand ration of 1), we consider two more scenarios where the supply-demand ratio is equal to 1.10 in one and 0.96 in the other. The results of these experiments are summarized in Figure 4.2.

First, we observe that in the case where the overall supply is smaller than the overall demand (i.e., $S/D = 0.96$), the difference in terms of the average age of blood transfused between the FIFO and LIFO issuing policies is not significant. This happens because the hospital does not have much choice in the way it issues the blood units due to the scarcity of the supply. In other words, for most days, the hospital should use all the blood on hand. Consequently, we see that the units do not stay in the inventory for long periods and even under the FIFO policy, the average age of blood transfused is 30 days. For the same reason, we cannot gain much improvement by issuing blood according to our ADP-based policy.

In contrast, when the overall supply is more than the demand (i.e., $S/D = 1.10$), the choice of the issuing policy is very important. In such a system, the hospital inevitably faces outdates, but having these outdates do not necessarily results in having shortages because of the excess supply. Consequently, we can issue fresher blood and still keep the shortage proportion very low. We observe that for the same shortage rate, the ADP-based policy achieves a significantly smaller average age of blood transfused. As we mentioned earlier, this is due to the fact that ADP-based policy only

Figure 4.2: Comparison of ADP-based policies (solid lines) and age-based threshold policies (dashed lines) for different supply-demand ratios.

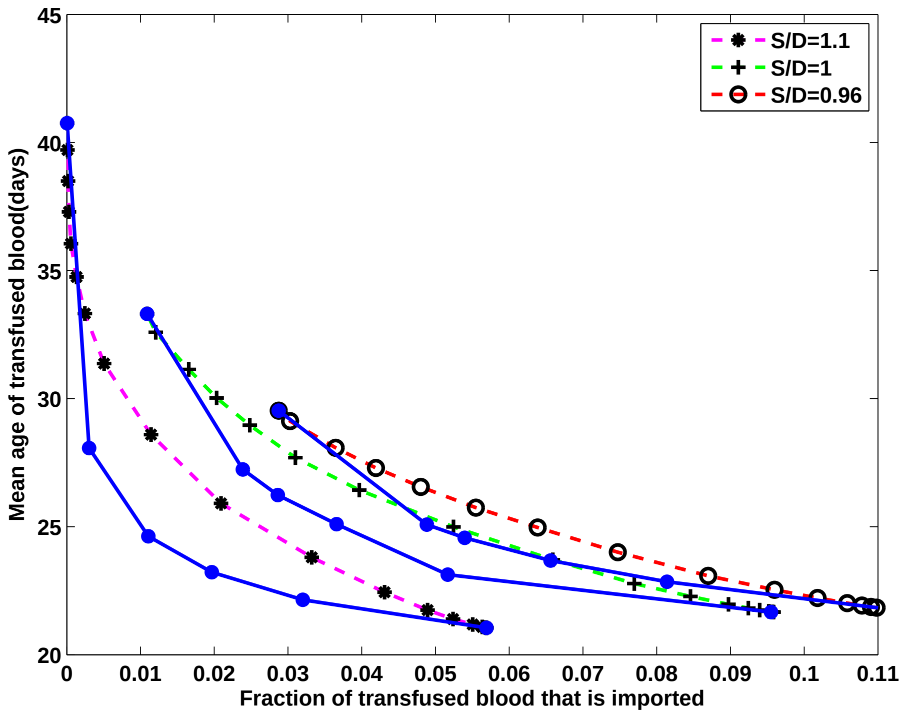

issues older units if the current inventory suggests a considerable chance of a shortage in the near future. In summary, our results suggest that using ADP-based policies is more beneficial when there is more flexibility in how we can issue blood to satisfy the demand.

Simple threshold policy inspired by ADP-based policy

In practice, using the ADP-based policy requires solving the problem presented in (4.15) in each period to determine which blood units should be issued to meet the demand. Aside from requiring software and an interface to do this, the solution might not have nice structure. This may cause concern among users and may limit the uptake of our results. Therefore, we studied the form of our ADP-based issuing policies to see if they might be translated into a simpler class of policies, analogous to the age-based threshold policies proposed by Atkinson et al. [7].

Our experiments based on the ADP-based issuing policies revealed several interesting observations. First, we observed that in cases where the remaining inventory at the end of the period is relatively low and a shortage situation in the next period is more likely, the ADP-based policy issues blood according to the FIFO rule. Second, the solution of the ADP often has multiple optimal solutions, including the issuance of blood in consecutive order of age. Using these observations, we proposed the following quantity-based threshold policy: for a fixed threshold κ, issue blood in consecutive order starting from units older than the freshest κ units. If the blood older than the first κ units is not sufficient to meet the demand, then issue blood from the κ^{th} oldest unit and younger (i.e., FIFO issuing policy).

Similar to the ADP-based policy, the threshold policy proposed above also issues blood in consecutive order of age and on a FIFO basis when the current inventory is low (i.e., when the blood older than the first κ units is not sufficient to meet the demand). However, this threshold policy is static because a fixed threshold κ is used regardless of the current state of the inventory. In contrast to the threshold policies proposed in Atkinson et al. [7], the threshold in our policy is defined for the quantity of blood units, and not the age of blood.

We investigated the performance of our quantity-based threshold policy using the simulation program by trying different values for the threshold κ. The results of our experiments are presented in Figure 4.3 for values of $\kappa = 50, 100, \ldots, 500$. We observe in Figure 4.3 that our quantity-based threshold policy outperforms the age-based threshold policy and performs nearly as good as ADP-based policy.

The quantity-based threshold policy performs better because it explicitly considers the total remaining inventory at the end of period in making issuing decisions whereas the age-based threshold policy ignores any information regarding the quantity of blood units of different ages. Furthermore, the age-based policy let the blood of age 42 days expire and only issues these units if there are no fresher blood units available. However, if the current overall inventory suggests that a shortage is likely in future, the quantity-based threshold policy issues the oldest blood and keeps the fresher blood for future use.

The quantity-based policy is easier to characterize and implement in practice compared to the ADP-based policy. By choosing the right value for a single parameter κ, one can achieve the right balance between the age of blood transfused and the shortage rate. As discussed above, the resulting policy also performs very close to the more complicated ADP-based policy.

Figure 4.3: Performance of quantity-based threshold policies for values of $\kappa = 50, 100, \ldots, 500$.

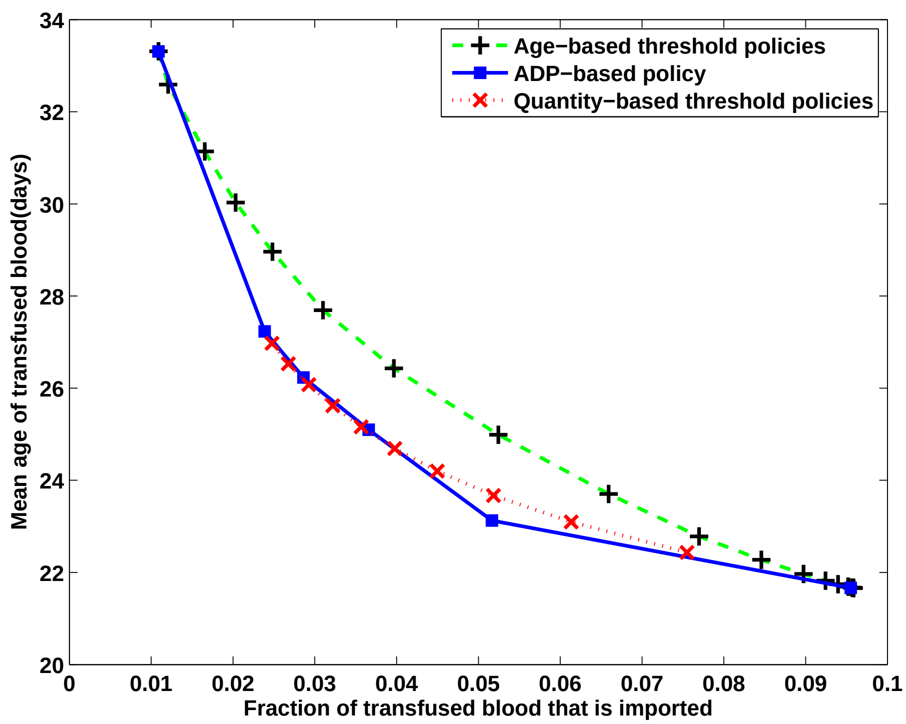

4.6 Conclusions

In this chapter, we studied the problem of issuing blood units in a hospital for transfusion when in addition to reducing shortages and outdates, the policy maker wants to decrease the average age of blood transfused. The optimal issuing policy in this case is not trivial and the dynamic programming formulation of this problem suffers from the curse of dimensionality. We overcome this difficulty by using ADP techniques. In particular, we define a set of basis functions to approximate the value function of the dynamic programming and solve the linear programming form of the dynamic programming using a column generation algorithm.

Our numerical results suggests that our ADP-based issuing policy (which is state-dependent) outperforms other classes of policies previously suggested in the literature. The ADP-based policy performs better than the threshold policy because it considers the on-hand inventory vector in making issuing decisions and adapts itself to the current state of the system, whereas the threshold policy is static and always use the same issuing rule based on a fixed pre-chosen threshold. Our results are especially of interest to policy-makers as it shows the average age of blood transfused can be significantly reduced at the expense of a reasonable increase in the shortage rate. Finally, we showed that there is more value in using state-dependent policies when the supply is greater than the demand.

Our work in this chapter also opens up several interesting directions for further research. The first natural extension to this work is considering different blood types and the possibility of issuing compatible but nonidentical blood units to satisfy the demand. Another important extension to this work is finding the optimal ordering policies assuming that the hospital issues blood according to the ADP-based issuing policy proposed in this chapter. More interestingly, one can jointly optimize the ordering and issuing policies. Finally, the blood supplier (e.g., Canadian Blood Services) also allocates blood from their inventory to satisfy the demand of different hospitals. Considering the fact that the average age of blood transfused is important, one might consider developing a model for finding the optimal allocation of blood of different ages to different hospitals.

Chapter 5

Conclusion, Extensions and Further Application

The research in this dissertation focused on three different decision making problems with applications of the stochastic optimization models in health care as a central theme. In this section, we provide a review of these problems, our solution approaches, and the main results. Furthermore, for each problem, we discuss possible extensions and avenues for further research.

First, in Chapter 2, we provided a novel framework for finding good screening policies for patients on the kidney transplant waiting list. We developed an analytical model for the single-patient case, and proved several properties of the optimal screening policy. Then, we used these properties and developed an efficient binary search algorithm to obtain the optimal policy. Finally, we incorporated our solution for the single-patient screening problem into a discrete event simulation model that heuristically (and dynamically) updates the screening policy for each patient on the waiting list. We also performed several numerical experiments using real data and made several important observations regarding good screening policies.

Analytically, we were able to show that the screening intervals should be decreasing under reasonable assumptions. This property contrasts sharply with the current practice of screening patients over fixed intervals where the interval only depends on the patients' risk for developing cardiovascular disease. Our numerical results suggested that screening guidelines should consider patients' remaining waiting time (or their rank on the waiting list) as a main factor in designing screening guidelines.

In practice, our results can help policy makers in designing better screening guidelines. In particular, as we showed, considering factors that affect the waiting time of patients in developing screening guideline could not only save money, but significantly reduce the likelihood of offering a kidney to a patient with unknown cardiovascular disease.

Since our analytical framework does not consider all the complexities of the real problem, we suggested heuristic methods to include several issues faced in practice. Alternatively, one can use other approximation methods such as approximate dynamic programming (ADP) to model and solve this problem. The main component of an ADP model is a set of basis functions for approximating the value function, for which, the insights and results from our work can be found useful.

In Chapter 3, we extended our previous results and developed a model to find the right timing for the inspections and replacements of a component that fails silently and is needed at a random future time. We developed an algorithm to obtain the optimal inspection policy that consists of scheduling a finite number of inspections as well as an age-based preventive replacement scheduled after the last inspection.

As directions for future research, one might consider the case where the inspections are not error-free. Furthermore, it is interesting to study the case where inspections and replacements take non-negligible time to perform. Finally, one might consider extending our results to the case where the back-up component replaces a main unit and the inspection/replacement decisions of both components need to be considered.

Finally, in Chapter 4 we examined the problem of finding the optimal issuing policy for the blood units in a hospital. Our modeling framework tries to balance the existing trade-off between minimizing shortages and minimizing the average age of blood transfused. We used approximate dynamic programming methods to design an algorithm that finds state-dependent approximate optimal issuing policies and showed that efficiencies can be achieved by using our ADP-based issuing policies compared to simpler static policies currently being used in practice or previously proposed in the literature.

Our results are useful for policy makers by providing evidence that the age of transfused blood can be reduced without a significant increase in the rate of shortages. Furthermore, we introduce simple to use issuing policies that only use the oldest blood when a shortage scenario in the near future is very likely. Using these policies, better health outcomes can be achieved while the current standards of efficiency are also met.

Our work in Chapter 4 can be extended in several ways. We considered only a single blood type in our analysis. However, in practice compatible (but nonidentical) blood units can be issued to satisfy the demand. It would be interesting to extend our results to the case where we manage the inventory of all the different blood types. Another important problem that hospital blood banks face is determining how many units of blood to order in each period. It is obvious that the ordering decision can affect the issuing policy and vice versa. Thus, one might consider optimizing both decision simultaneously. We believe that our work provides the first step towards this goal and insights from our results might prove valuable when we consider this more complicated problem. Finally, at the system level, it is interesting to study the problem of how the supplier should allocate the blood units to different hospitals, given the new issuing policies suggested in this dissertation.

Bibliography

[1] H. Abouee-Mehrizi, O. Baron, O. Berman, and V. Sarhangian. Allocation policies in blood transfusion. Working paper, 2013.

[2] M. Ahrens. Smoke alarms in u.s. home fires. Technical report, National Fire Protection Association, Fire Analysis and Research Division, Sep 2011.

[3] O. Alagoz, L. M. Maillart, A. J. Schaefer, and M. S. Roberts. The optimal timing of living-donor liver transplantation. *Management Science*, 50(10):1420–1430, 2004.

[4] O. Alagoz, L. M. Maillart, A. J. Schaefer, and M. S. Roberts. Choosing among living-donor and cadaveric livers. *Management Science*, 53(11): 1702–1715, 2007.

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

[6] K. M. Anderson, P. M. Odell, P. W.F. Wilson, and W. B. Kannel. Cardiovascular disease risk profiles. *American Heart Journal*, 121(1, Part 2):293 – 298, 1991.

[7] M. P. Atkinson, M. J. Fontaine, L. T. Goodnough, and L. M. Wein. A novel allocation strategy for blood transfusions: investigating the trade-off between the age and availability of transfused blood. *Transfusion*, 52(1):108–117, 2012.

[8] T. Aven and B. Bergman. Optimal replacement times: A general set-up. *Journal of Applied Probability*, 23(2):432 – 442, 1986.

[9] T. Ayer, O. Alagoz, and N. K. Stout. A POMDP approach to personalize mammography screening decisions. *Operations Research*, 2012. (to appear).

[10] R. E. Barlow and F. Proschan. *Mathematical Theory of Reliability*. Wiley, New York, 1965.

[11] R. E. Barlow, L. C. Hunter, and F. Proschan. Optimum checking procedures. *J. Soc. Indust. Appl. Math.*, 11(4):1078–1095, 1963.

[12] B.C. Transplant Stats, 2011. URL `http://www.transplant.bc.ca/stats_faqs_main.htm`. accessed Dec 25, 2011.

[13] D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1996.

[14] D. Bertsimas, V. F. Farias, and N. Trichakis. On the efficiency-fairness trade-off. *Management Science*, 2012.

[15] A. Chelbi and D. Ait-Kadi. Inspection strategies for randomly failing systems. In M. Ben-Daya, S. O. Duffuaa, A. Raouf, J. Knezevic, and D. Ait-Kadi, editors, *Handbook of Maintenance Management and Engineering*. Springer London, 2009.

[16] W. L. Cooper. Pathwise properties and performance bounds for a perishable inventory system. *Operations Research*, 49(3):455–466, 2001.

[17] G. M. Danovitch, S. Hariharan, J. D. Pirsch, D. Rush, D. Roth, E. Ramos, R. C. Starling, C. Cangro, and M. R. Weir. Management of the waiting list for cadaveric kidney transplants: Report of a survey and recommendations by the clinical practice guidelines committee of the american society of transplantation. *Journal of the American Society of Nephrology*, 13:528–535, 2002.

[18] G. M. Danovitch, S. Hariharan, J. D. Pirsch, D. Rush, D. Roth, E. Ramos, R. C. Starling, C. Cangro, and M. R. Weir. Management of the waiting list for cadaveric kidney transplants: Report of a survey and recommendations by the clinical practice guidelines committee of the american society of transplantation. *Journal of the American Society of Nephrology*, 13:528–535, 2002.

[19] B. Deniz, I. Karaesmen, and A. Scheller-Wolf. Managing perishables with substitution: inventory issuance and replenishment heuristics. *Manufacturing & Service Operations Management*, 12(2):319–329, 2010.

[20] L. Dieulle, C. Brenguer, A. Grall, and M. Roussignol. Sequential condition-based maintenance scheduling for a deteriorating system. *European Journal of Operational Research*, 150(2):451 – 461, 2003.

[21] R. N. Foley, P. S. Parfrey, and M. J. Sarnak. Clinical epidemiology of cardiovascular disease in chronic renal disease. *American Journal of Kidney Diseases*, 32(5, Supplement 3):S112 – S119, 1998.

[22] B. E. Fries. Optimal ordering policy for a perishable commodity with fixed lifetime. *Operations Research*, 23(1):46–61, 1975.

[23] R. S. Gaston, G. M. Danovitch, P. L. Adams, J. J. Wynn, R. M. Merion, M. H. Deierhoi, R. A. Metzger, J. M. Cecka, W. E. Harmon, A. B. Leichtman, A. Spital, E. Blumberg, C. A. Herzog, R. A. Wolfe, D. B. Tyan, J. Roberts, R. Rohrer, F. K. Port, and F. L. Delmonico. The report of a national conference on the wait list for kidney transplantation. *American Journal of Transplantation*, 3(7):775–785, 2003.

[24] J. S. Gill, I. Ma, D. Landsberg, N. Johnson, and A. Levin. Cardiovascular events and investigation in patients who are awaiting cadaveric kidney transplantation. *Journal of the American Society of Nephrology*, 16(3):808–816, 2005.

[25] R. Haijema, J. van der Wal, and N. M. van Dijk. Blood platelet production: Optimization by dynamic programming and simulation. *Computers & Operations Research*, 34(3):760–779, 2007.

[26] A. Humar, S. R. Kerr, T. Ramcharan, K. J. Gillingham, and A. J. Matas. Peri-operative cardiac morbidity in kidney transplant recipients: incidence and risk factors. *Clinical Transplantation*, 15(3):154 – 158, 2001.

[27] A. K. S. Jardine and J. A. Buzacott. Equipment reliability and maintenance. *European Journal of Operational Research*, 19(3):285 – 296, 1985.

[28] I. Z. Karaesmen, A. Scheller-Wolf, and B. Deniz. Managing perishable and aging inventories: review and future research directions. In *Planning production and inventories in the extended enterprise*, pages 393–436. Springer, 2011.

[29] H. Kaspi and D. Perry. Inventory systems of perishable commodities. *Advances in Applied Probability*, pages 674–685, 1983.

[30] H. Kaspi and D. Perry. Inventory systems for perishable commodities with renewal input and poisson output. *Advances in Applied Probability*, pages 402–421, 1984.

[31] Y. H. Kim and L. C. Thomas. Repair strategies in an uncertain environment: Markov decision process approach. *The Journal of the Operational Research Society*, 57(8):957–964, 2006.

[32] R. L. A. Kirch and M. Klein. Surveillance schedules for medical examinations. *Management Science*, 20(10):1403–1409, 1974.

[33] C. G. Koch, L. Li, D. I. Sessler, P. Figueroa, G. A. Hoeltge, T. Mihaljevic, and E. H. Blackstone. Duration of red-cell storage and complications after cardiac surgery. *New England Journal of Medicine*, 358(12):1229–1239, 2008.

[34] H. Luss and Z. Kander. Inspection policies when duration of checkings is non-negligible. *Operational Research Quarterly*, 25(2):299–309, 1974.

[35] L. M. Maillart, J. S. Ivy, S. Ransom, and K. Diehl. Assessing dynamic breast cancer screening policies. *Operations research*, 56(6):1411–1427, 2008.

[36] L.M. Maillart and S.M. Pollock. Cost-optimal condition-monitoring for predictive maintenance of 2-phase systems. *IEEE Transactions on Reliability*, 51(3):322–330, 2002.

[37] J. J. McCall. Maintenance policies for stochastically failing equipment: A survey. *Management Science*, 11(5):493–524, 1965.

[38] P. A. McCullough. Coronary artery disease. *Clinical Journal of the American Society of Nephrology*, 2(3):611 – 616, 2007.

[39] A. G. Munford and A. K. Shahani. A nearly optimal inspection policy. *Operational Research Quarterly (1970-1977)*, 23(3):pp. 373–379, 1972.

[40] S. Nahmias. Optimal ordering policies for perishable inventory–ii. *Operations Research*, 23(4):735–749, 1975.

[41] S. Nahmias. Myopic approximations for the perishable inventory problem. *Management Science*, 22(9):1002–1008, 1976.

[42] S. Nahmias. Higher-order approximations for the perishable-inventory problem. *Operations Research*, 25(4):630–640, 1977.

[43] S. Nahmias. Perishable inventory theory: A review. *Operations Research*, 30(4):680–708, 1982.

[44] S. Nahmias. *Perishable Inventory Systems*, volume 160. Springer, 2011.

[45] S. Nahmias and W. P. Pierskalla. Optimal ordering policies for a product that perishes in two periods subject to stochastic demand. *Naval Research Logistics Quarterly*, 20(2):207–229, 1973.

[46] T. Nakagawa. Optimum inspection policies for a standby unit. *Journal of the Operations Research Society of Japan*, 23:13–26, 1980.

[47] T. Nakagawa. Maintenance and optimum policy. In Hoang Pham, editor, *Handbook of Reliability Engineering*. Springer, London, 2003.

[48] National Kidney Foundation. K/doqi clinical practice guidelines for cardiovascular disease in dialysis patients. *American Journal of Kidney Diseases*, 45(4):S1–S153, 2005.

[49] P. M. Odell, K. M. Anderson, and W. B. Kannel. New models for predicting cardiovascular events. *Journal of Clinical Epidemiology*, 47 (6):583 – 592, 1994.

[50] A. O. Ojo, J. A. Hanson, R. A. Wolfe, A. B. Leichtman, L. Y. Agodoa, and F. K. Port. Long–term survival in renal transplant recipients with graft function. *Kidney International*, 57:307–313, 2000.

[51] OPTN. Organ procurement and transplantation network (OPTN) national data report: kidney kaplan-meir median waiting times for registrations listed 1999–2004, 2011. URL `http://optn.transplant.hrsa.gov/latestData/viewDataReports.asp`. accessed Dec 25, 2011.

[52] P. S. Parfrey and R. N. Foley. The clinical epidemiology of cardiac disease in chronic renal failure. *Journal of the American Society of Nephrology*, 10(7):1606–1615, 1999.

[53] M. Parlar, D. Perry, and W. Stadje. Fifo versus lifo issuing policies for stochastic perishable inventory systems. *Methodology and Computing in Applied Probability*, 13(2):405–417, 2011.

[54] G. Parmigiani. Inspection times for stand-by units. *Journal of Applied Probability*, 31(4):1015–1025, 1994.

[55] G. Parmigiani. Optimal scheduling of fallible inspections. *Operations Research*, 44(2):360–367, 1996.

[56] R. E. Patterson, R.t L. Eisner, and S. F. Horowitz. Comparison of cost-effectiveness and utility of exercise ECG, single photon emission computed tomography, positron emission tomography, and coronary angiography for diagnosis of coronary artery disease. *Circulation*, 91 (1):54–65, 1995.

[57] D. Perry. Analysis of a sampling control scheme for a perishable inventory system. *Operations Research*, 47(6):966–973, 1999.

[58] W. P. Pierskalla and C. D. Roach. Optimal issuing policies for perishable inventory. *Management Science*, 18(11):603 – 614, 1972.

[59] 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.

[60] F. K. Port, R. A. Wolfe, E. A. Mauger, D. P. Berling, and K. Jiang. Comparison of survival probabilities for dialysis patients vs cadaveric renal transplant recipients. *JAMA: The Journal of the American Medical Association*, 270(11):1339–1343, 1993.

[61] M. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. Wiley-Interscience, New York, 1994.

[62] A. Sabouri, W. T. Huh, and S. M. Shechter. Screening strategies for patients on the kidney transplant waiting list. 2013. (working paper).

[63] A. K. Salahudeen, N. Haider, and W. May. Cold ischemia and the reduced long-term survival of cadaveric renal allografts. *Kidney International*, 65:713–718, 2004.

[64] B. Sandikci, L. M. Maillart, A. J. Schaefer, O. Alagoz, and M. S. Roberts. Estimating the patients price of privacy in liver transplantation. *Operations Research*, 56(6):1393–1410, 2008.

[65] M. J. Sarnak and A. S. Levey. Epidemiology of cardiac disease in dialysis patients. *Seminars in Dialysis*, 12(2):69 – 76, 1999.

[66] P. Schnuelle, D. Lorenz, M. Trede, and F. J. Van Der Woude. Impact of renal cadaveric transplantation on survival in end-stage renal failure: evidence for reduced mortality risk compared with hemodialysis during long-term follow-up. *Journal of the American Society of Nephrology*, 9 (11):2135–2141, 1998.

[67] P. J. Schweitzer and A. Seidmann. Generalized polynomial approximations in markovian decision processes. *Journal of mathematical analysis and applications*, 110(2):568–582, 1985.

[68] B. Sengupta. Inspection procedures when failure symptoms are delayed. *Operations Research*, 28(3):768–776, 1980.

[69] Y. S. Sherif and M. L. Smith. Optimal maintenance models for systems subject to failure–a review. *Naval Research Logistics Quarterly*, 28(1): 47–74, 1981.

[70] M. Shwartz. A mathematical model used to analyze breast cancer screening strategies. *Operations Research*, 26(6):937–955, 1978.

[71] X. Su and S. A. Zenios. Patient choice in kidney allocation: A sequential stochastic assignment model. *Operations Research*, 53(3):443–455, 2005.

[72] X. Su and S. A. Zenios. Recipient choice can address the efficiency-equity trade-off in kidney transplantation: A mechanism design model. *Management Science*, 52(11):1647–1660, 2006.

[73] L. C. Thomas, P. A. Jacobs, and P. Gaver. Optimal inspection policies for standby systems. *Communications in Statistics-Stochastic Models*, 3(2):259–273, 1987.

[74] USRDS. *Annual Data Report: 2011 Atlas of CKD and ESRD*. United States Renal Data System, 2011.

[75] C. Valdez-Flores and R. M. Feldman. A survey of preventive maintenance models for stochastically deteriorating single-unit systems. *Naval Research Logistics*, 36(4):419–446, 1989.

[76] D. Wang, J. Sun, S. B. Solomon, H. G. Klein, and C. Natanson. Transfusion of older stored blood and risk of death: a meta-analysis. *Transfusion*, 52(6):1184 – 1195, 2012.

[77] H. Wang. A survey of maintenance policies of deteriorating systems. *European Journal of Operational Research*, 139(3):469 – 489, 2002.

[78] L. Yeh. An optimal inspection-repair-replacement policy for standby systems. *Journal of Applied Probability*, 32(1):212–223, 1995.

[79] G. Zallen, P. J. Offner, E. E. Moore, J. Blackwell, D. J. Ciesla, J. Gabriel, C. Denny, and C. C. Silliman. Age of transfused blood is an independent risk factor for postinjury multiple organ failure. *The American Journal of Surgery*, 178(6):570 – 572, 1999.

[80] S. A. Zenios. Optimal control of a paired-kidney exchange program. *Management Science*, 48(3):328–342, 2002.

[81] 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.

[82] S. Zhang, P. Hanagal, P. I. Frazier, A. J. Meltzer, and Schneider D. B. Optimal patient-specific post-operative surveillance for vascular surgery. *Proceedings of the 7th INFORMS Workshop on Data Mining and Health Informatics (DM-HI 2012)*, 2012.

Appendix A

Chapter 2: Supporting Results

First we define function $\Omega(t, x, y)$, which will be used in Lemma 3 to characterize the first order necessary condition. Furthermore, we prove several properties of this function and use them to establish properties of the optimal screening policy.

$$\Omega(t, x, y) := \frac{F(t) - F(t-x)}{f(t)} - c_p \int_0^y \frac{g(t+u)\bar{H}(t+u)}{\Lambda(t)g(t)\bar{H}(t)} du \\ - \frac{c_s}{\Lambda(t)} \Big[\frac{\bar{F}(t)}{f(t)} \frac{c_p - \Lambda(t)}{c_s} + \frac{\bar{G}(t+y)\bar{H}(t+y)}{g(t)\bar{H}(t)} \Big], \tag{A.1}$$

where

$$\Lambda(t) = c_p - c_s \big(1 + \frac{h(t)\bar{G}(t)}{g(t)\bar{H}(t)}\big), \tag{A.2}$$

is the function introduced in Assumption 1. We remark that $\Lambda(t) < c_p$ since the expression in the parentheses is positive.

Lemma 3 *Any optimal screening policy $\mathcal{T}^* = \{t_j^*\}_{j=1}^n$ satisfies $0 < t_1^* < t_2^* < \ldots$. Furthermore, a necessary condition for optimality of a screening policy is*

$$\Omega(t_j^*, x_j^*, x_{j+1}^*) = 0 \qquad \textit{for each } 1 \leq j. \tag{A.3}$$

Proof. Recall that the objective is to minimize $\mathcal{C}(\mathcal{T})$ subject to $0 \leq t_1 \leq t_2 \leq \ldots$. We need to show that the optimal screening policy cannot be a boundary solution. Suppose, by way of contradiction, that a screening policy, $\mathcal{T}$, satisfies $t_m = t_{m-1}$ for some $m \geq 1$. We will show that this screening policy cannot be optimal. Let $\mathcal{T}' = \{t'_j\}_{j=1}^\infty$ be a new screening policy in which, $t'_j = t_j$ for $j < m$ and $t'_j = t_{j+1}$ for $j \geq m$. In other words, $\mathcal{T}'$ performs a screening at every time $\mathcal{T}$ does, but at time t_m, $\mathcal{T}'$

performs one less screening than $\mathcal{T}$. Using (2.1), it can be shown that for the new screening policy, the expected penalty cost remains unchanged while it removes the redundancy of screenings at time t_m. More specifically, the expected screening cost of the new screening policy is smaller by the product of c_s and the probability that the patient is alive, CVD-free and still on the waiting list by time t_m, i.e., $c_s\bar{F}(t_m)\bar{G}(t_m)\bar{H}(t_m)$. Therefore, $\mathcal{T}'$ achieves lower expected total cost, from which we conclude that the optimal screening policy is an interior solution. Thus, it should satisfy the first order necessary condition: $\frac{\partial \mathcal{C}(\mathcal{T})}{\partial t_j} = 0$ for each $j \geq 1$. By differentiating the expression of $\mathcal{C}(\mathcal{T})$ with respect to t_j and setting it equal to zero, after simplification and rearrangement, we obtain for each $j \geq 1$:

$$\begin{aligned} &\Lambda(t_j)\frac{F(t_j) - F(t_{j-1})}{f(t_j)} - c_p \int_{t_j}^{t_{j+1}} \frac{g(u)\bar{H}(u)}{g(t_j)\bar{H}(t_j)} du \\ &- c_s \Big[\frac{\bar{F}(t_j)}{f(t_j)}\frac{c_p - \Lambda(t_j)}{c_s} + \frac{\bar{G}(t_{j+1})\bar{H}(t_{j+1})}{g(t_j)\bar{H}(t_j)}\Big] = 0, \end{aligned} \tag{A.4}$$

where, $\Lambda(t)$ is defined in (A.2). Note that since $\Lambda(t_j) < c_p$, (A.4) implies $\Lambda(t_j) > 0$; otherwise, the left hand side of (A.4) is negative. After dividing both sides by $\Lambda(t_j)$, the left hand side is equal to $\Omega(t_j, x_j, x_{j+1})$. □

Lemma 4 *Suppose both the kidney arrival time and the death time follow Weibull distribution, i.e.,* $g(t) = \frac{\alpha}{\lambda}(\frac{t}{\lambda})^{\alpha-1}e^{-(\frac{t}{\lambda})^{\alpha}}$ *and* $h(t) = \frac{\beta}{\mu}(\frac{t}{\mu})^{\beta-1}e^{-(\frac{t}{\mu})^{\beta}}$. *Also, assume that* $\alpha > \beta \geq 1$. *Then, the following statements hold for function* $\Lambda(t)$ *defined in (A.2):*

(a) $\Lambda(t)$ *is increasing in* t.

(b) There exists some $\hat{\psi}$ *such that* $\Lambda(t)\bar{H}(t)g(t)$ *is strictly decreasing in* t *for all* $t > \hat{\psi}$.

Proof. (a) This statement holds for $\alpha > \beta$ since $\frac{h(t)\bar{G}(t)}{g(t)\bar{H}(t)} = \frac{\beta\lambda^{\alpha}}{\alpha\mu^{\beta}}t^{\beta-\alpha}$ is decreasing in t.
(b) By differentiation and after simplification, we can write:

$$(\Lambda(t)\bar{H}(t)g(t))' = e^{-(\frac{t}{\lambda})^{\alpha}}e^{-(\frac{t}{\mu})^{\beta}}\Big[-(c_p - c_1 - c_2)\frac{\alpha^2}{\lambda^{2\alpha}}t^{2(\alpha-1)} + \mathcal{O}(t^k)\Big],$$

where $k < 2(\alpha - 1)$. Then, the square bracket approaches $-\infty$ as $t \to +\infty$. It follows that there exists some $\hat{\psi}$ such that $(\Lambda(t)\bar{H}(t)g(t))'$ is negative for all $t > \hat{\psi}$. □

Lemma 5 *For PF_2 densities f and g, and for $(t,x,y) \in \mathbb{R}_+^3$ such that $\Lambda(t) > 0$, the following statements hold:*

(a) $\Omega(t,x,y)$ is increasing in t.

(b) $\Omega(t,x,y)$ is strictly increasing in x for all y and for all t such that $t-x>0$.

(c) $\Omega(t,x,y)$ is strictly decreasing in y for all t and x.

Proof. (a) The following facts suffice to prove that $\Omega(t,x,y)$ is increasing in t: $\frac{F(t)-F(t-x)}{f(t)}$ is increasing in t and $\frac{\bar{F}(t)}{f(t)}$ is decreasing in t (Theorem 3 and Corollary 3.1 in Barlow et al. [11]); $\Lambda(t)$ is increasing in t (Assumption 1(a)) and $\Lambda(t) < c_p$; $\frac{g(t+u)}{g(t)}$ is decreasing in t for $u > 0$ since density g is PF_2 (see Barlow et al. [11]); and $\frac{\bar{H}(t+u)}{\bar{H}(t)} = 1 - \frac{H(t+u)-H(t)}{\bar{H}(t)}$ is decreasing in t since density h has IFR property.
(b) Since the density function f satisfies $f(a) > 0$ for any $a > 0$, we have $\frac{\partial \Omega(t,x,y)}{\partial x} = \frac{f(t-x)}{f(t)} > 0$. Thus, $\Omega(t,x,y)$ is strictly increasing in x.
(c) Since function $g(t)$ and $\bar{H}(t)$ are positive for any $t > 0$, we have $\frac{\partial \Omega(t,x,y)}{\partial y} = -\frac{\Lambda(t+y)\bar{H}(t+y)g(t+y)}{\Lambda(t)\bar{H}(t)g(t)} < 0$. Thus, $\Omega(t,x,y)$ is strictly decreasing in y. □

Lemma 6 *Let ψ_f be the mode of f. For $(t,x,y) \in \mathbb{R}_+^3$ such that $t-x > \max(\psi_f, \hat{\psi})$, where $\hat{\psi}$ is defined in Assumption 1(b), the following statements hold:*

(a) $\Omega(t,rx,ry) > r\Omega(t,x,y)$ for any $r > 1$.

(b) $\Omega(t,x+\delta,y+\delta) > \Omega(t,x,y)$ for any $\delta > 0$.

Proof. (a) Since $f'(t) < 0$ for $t > \psi_f$ (see Theorem 2 in Appendix 1 of Barlow and Proschan [10]) and $t-x > \psi_f$, we have $A := \frac{\partial^2 \Omega(t,x,y)}{\partial x^2} = \frac{-f'(t-x)}{f(t)} > 0$. Similarly, for $t > \hat{\psi}$, we have $B := \frac{\partial^2 \Omega(t,x,y)}{\partial y^2} = -\frac{(\Lambda(t+y)g(t+y)\bar{H}(t+y))'}{\Lambda(t)g(t)\bar{H}(t)} > 0$ by Assumption 1(b). It follows that for $t-x > \max(\psi_f, \hat{\psi})$, the Hessian matrix $\nabla^2 \Omega(t,x,y) = \begin{bmatrix} A & 0 \\ 0 & B \end{bmatrix}$ is positive definite (note we take t as constant), i.e., the function $\Omega(t,x,y)$ is strictly convex in (x,y) for all $t > x + \max(\psi_f, \hat{\psi})$. Now, for any $r > 1$, we can write:

$$\Omega(t,x,y) < \frac{1}{r}\Omega(t,rx,ry) + \frac{r-1}{r}\Omega(t,0,0) < \frac{1}{r}\Omega(t,rx,ry),$$

where the first inequality follows from the strict convexity of $\Omega(t,x,y)$ in (x,y), and the second inequality holds since $\Omega(t,0,0) < 0$ (see (A.1)).
(b) Since we have $f'(t) < 0$ for any $t > \psi_f$, it follows that $f(t-x) > f(t)$ for $t - x > \psi_f$. Therefore, $\frac{\partial\Omega(t,x,y)}{\partial x} = \frac{f(t-x)}{f(t)} > 1$ for all x and t such that $t - x > \psi_f$ and for all y. Similarly, one can show that for $t > \hat{\psi}$, we have $\frac{\partial\Omega(t,x,y)}{\partial y} = \frac{-\Lambda(t+y)g(t+y)\bar{H}(t+y)}{\Lambda(t)g(t)\bar{H}(t)} > -1$. Now, since $\Omega(t,x,y)$ is strictly convex in (x,y) for all $t > x + \max(\psi_f, \hat{\psi})$ (see proof of part (a)), for any $\delta > 0$, we can write

$$\Omega(t, x+\delta, y+\delta) > \Omega(t,x,y) + \frac{\partial\Omega(t,x,y)}{\partial x}\delta + \frac{\partial\Omega(t,x,y)}{\partial y}\delta > \Omega(t,x,y),$$

where the first inequality comes from the strict convexity of $\Omega(t,x,y)$ in (x,y), and the second inequality holds since $\frac{\partial\Omega(t,x,y)}{\partial x} > 1$ and $\frac{\partial\Omega(t,x,y)}{\partial y} > -1$.
□

Lemma 7 *Suppose that $\mathcal{T} = \{t_j\}_{j=1}^{\infty}$ obtained recursively from (A.3). If there exists $m \geq 1$ such that $x_{m+1} > x_m$, then $x_{j+1} > x_j$ for all $j \geq m$.*

Proof. It suffices to show that if $x_{m+1} > x_m$ for some m, then $x_{m+2} > x_{m+1}$ follows. From (A.3), we have

$$\Omega(t_{m+1}, x_{m+1}, x_{m+2}) = \Omega(t_m, x_m, x_{m+1}) = 0. \tag{A.5}$$

Note that $t_{m+1} > t_m$ and $x_{m+1} > x_m$. Now, since $\Omega(t,x,y)$ is increasing in t, strictly increasing in x and strictly decreasing in y (Lemma 5), we must have $x_{m+2} > x_{m+1}$ in order to satisfy (A.5). This completes the proof. □

Lemma 8 *For any $j \geq 1$, x_j and t_j obtained recursively from (A.3) are strictly increasing in t_1.*

Proof. Since $x_1 = t_1$, from (A.3) for $j = 1$, we have $\Omega(t_1, t_1, x_2) = 0$. To satisfy this equation, x_2 must be strictly increasing in t_1 since $\Omega(t,x,y)$ is increasing in t, strictly increasing in x and strictly decreasing in y (Lemma 5). Consequently, $t_2 = t_1 + x_2$ is strictly increasing in t_1. Now by inductive reasoning it follows that t_j and x_j are strictly increasing in t_1 for any $j \geq 1$.
□

Appendix B

Chapter 2: Proofs of Main Results

Proof of Theorem 1. Suppose, by way of a contradiction, that in an optimal screening policy, the lengths of the screening intervals are not decreasing (i.e., $x_m > x_{m-1}$ for some m). Then by Lemma 7, x_j is strictly increasing in j for all $j \geq m$. In particular, there exists k such that $x_k > rx_{k-1}$ for some $r > 1$ and $t_{k-2} = t_{k-1} - x_{k-1} > \max(\psi_f, \hat{\psi})$, where ψ_f and $\hat{\psi}$ are defined in Lemma 6 (and the existence of such t_{k-2} is ensured by assumption that $\lim_{j\to\infty} t_j = \infty$). Now, we can write

$$\begin{aligned}\Omega(t_k, x_k, rx_k) &> \Omega(t_k, rx_{k-1}, rx_k) \geq \Omega(t_{k-1}, rx_{k-1}, rx_k) \\ &> r\Omega(t_{k-1}, x_{k-1}, x_k) = 0 = \Omega(t_k, x_k, x_{k+1}),\end{aligned}$$

where the first two inequalities hold since $\Omega(t, x, y)$ is strictly increasing in x and increasing in t (Lemma 5), the third inequality holds by Lemma 6(a), and the last two equalities hold by Lemma 3. It follow that $x_{k+1} > rx_k$ since $\Omega(t, x, y)$ is strictly decreasing in y (Lemma 5(c)). Then, using induction, we can show that $\lim_{j\to\infty} x_j = \infty$.

From (A.1) and Lemma 3, $\Omega(t_j, x_j, x_{j+1}) = 0$ implies

$$\begin{aligned}\frac{F(t_j) - F(t_j - x_j)}{f(t_j)} &= c_p \int_0^{x_{j+1}} \frac{\bar{H}(t_j + u)g(t_j + u)}{\Lambda(t_j)\bar{H}(t_j)g(t_j)} du \\ &\quad + \frac{c_s}{\Lambda(t_j)} \Big[\frac{\bar{F}(t_j)}{f(t_j)} \frac{c_p - \Lambda(t_j)}{c_s} + \frac{\bar{G}(t_j + x_{j+1})\bar{H}(t_j + x_{j+1})}{g(t_j)\bar{H}(t_j)} \Big] \\ &\leq c_p \int_0^{+\infty} \frac{\bar{H}(t_j + u)g(t_j + u)}{\Lambda(t_j)\bar{H}(t_j)g(t_j)} du \\ &\quad + \frac{c_s}{\Lambda(t_j)} \Big[\frac{\bar{F}(t_j)}{f(t_j)} \frac{c_p - \Lambda(t_j)}{c_s} + \frac{\bar{G}(t_j)}{g(t_j)} \Big], \qquad \text{(B.1)}\end{aligned}$$

where the inequality holds since the integrand is nonnegative and $\bar{G}(t)\bar{H}(t)$ is decreasing in t. Note that the right hand side of (B.1) is bounded since non-negative functions $\frac{\bar{F}(t)}{f(t)}$, $\frac{\bar{G}(t)}{g(t)}$ and $\frac{\bar{H}(t+u)g(t+u)}{\Lambda(t)H(t)g(t)}$ are decreasing in t (see proof of Lemma 5(a)). However, since $\frac{F(t)-F(t-x)}{f(t)}$ is strictly convex and strictly increasing in x, and increasing in t (proofs are similar to their counterparts in Lemma 6(a) and Lemma 5(a)-(b)), it follows that $x_j \to \infty$ and $t_j \to \infty$ imply that the left hand side of (B.1) is unbounded. This contradiction shows that the lengths of the optimal inspection intervals must be decreasing. □

Proof of Theorem 2. First, we show the "if part" of the both statements. Assume $t_1 > t_1^*$. Then by Lemma 8, $t_j > t_j^*$ and $x_j > x_j^* > 0$ for all j. Let $\delta_j = x_j - x_j^*$. Then, $\delta_j > 0$ for all j. Moreover, since $\lim_{j\to\infty} t_j^* = \infty$, there exists some k such that $t_{k-1}^* = t_k^* - x_k^* > \max(\psi_f, \hat{\psi})$ (where ψ_f and $\hat{\psi}$ are defined in Lemma 6). First, we establish that δ_j is strictly increasing for all $j \geq k$. We can write:

$$\begin{aligned}
\Omega(t_k, x_k^* + \delta_k, x_{k+1}^* + \delta_k) &\geq \Omega(t_k^*, x_k^* + \delta_k, x_{k+1}^* + \delta_k) \\
&> \Omega(t_k^*, x_k^*, x_{k+1}^*) \\
&= 0 \\
&= \Omega(t_k, x_k, x_{k+1}) \\
&= \Omega(t_k, x_k^* + \delta_k, x_{k+1}^* + \delta_{k+1}),
\end{aligned}$$

where the first inequality holds since $\Omega(t, x, y)$ is increasing in t and $t_k > t_k^*$, the second inequality follows from Lemma 6(b), and the last equality follows from the definition of δ_k and δ_{k+1} (the other two equalities hold by (A.3)). It follow that $\delta_{k+1} > \delta_k$ since $\Omega(t, x, y)$ is strictly decreasing in y (Lemma 5(c)). Now, we claim that $\lim_{k\to\infty} \delta_k = +\infty$. To prove this claim suppose, by way of contradiction, that $\lim_{k\to\infty} \delta_k = \hat{\delta}$. Then, we have:

$$\begin{aligned}
\lim_{k\to\infty} \Omega(t_k, x_k, x_{k+1}) &= \lim_{k\to\infty} \Omega(t_k, x_k^* + \delta_k, x_{k+1}^* + \delta_{k+1}) \\
&= \lim_{k\to\infty} \Omega(t_k, x_k^* + \hat{\delta}, x_{k+1}^* + \hat{\delta}) \\
&> \lim_{k\to\infty} \Omega(t_k, x_k^*, x_{k+1}^*) \\
&\geq \lim_{k\to\infty} \Omega(t_k^*, x_k^*, x_{k+1}^*) \\
&= 0,
\end{aligned}$$

where the first equality follows from the definition of δ_k and δ_{k+1}, the second equality holds since $\lim_{k\to\infty} \delta_k = \hat{\delta}$, the last equality holds by (A.3), and the inequalities hold by Lemma 6(b) and Lemma 5(a), respectively. However, $\Omega(t_k, x_k, x_{k+1}) = 0$ for all k and therefore cannot converge to a positive number. This contradiction shows that $\lim_{k\to\infty} \delta_k = +\infty$. It follows that there exists some m such that $\delta_m > t_1 = x_1$. Then, $x_m = x_m^* + \delta_m > x_1$. Then, we have $x_j > x_{j-1}$ for all $j \geq m$ by Lemma 7. This completes the "if part" of (A).

Now we consider the case that $t_1 < t_1^*$. Here, we have $\delta_j < 0$ (i.e., $t_j < t_j^*$ and $x_j < x_j^*$ for all j), and the rest of the proof can be derived similarly to show the existence of m such that $x_m < 0$. More specifically, all inequalities would be in the reverse direction, and we conclude that $\lim_{k\to\infty} \delta_k = -\infty$. Thus, there exists some m such that $\delta_m < -t_1^* = -x_1^*$. Then, $x_m = x_m^* + \delta_m < 0$ since $x_m^* < x_1^*$ by Theorem 1. Since both x_j^* and δ_j decrease as $j \geq m$ increases, we have $x_j < 0$ for all $j \geq m$. This completes the "if part" of (B).

To show the "only if part" of (A), suppose there exists some m such that $x_j > x_{j-1} > 0$ for all $j \geq m$, but $t_1 > t_1^*$ is not true, i.e., $t_1 \leq t_1^*$. For $t_1 = t_1^*$, we cannot have $x_m > x_{m-1}$ for any m by Theorem 1. Furthermore, $t_1 < t_1^*$ implies that there exists some q such that $x_j < 0$ for all $j \geq q$ by the "if part" of (B). This result contradicts our initial assumption. Thus, we must have $t_1 > t_1^*$.

Finally, we show that the "only if part" of (B) is also true. Suppose there exists some m such that $x_j < 0$ for all $j \geq m$, but $t_1 < t_1^*$ is not true, i.e., $t_1 \geq t_1^*$. Then, by Lemma 8, we have $x_j \geq x_j^* > 0$. This result contradicts our assumption that $x_m < 0$. It follows that $t_1 < t_1^*$. $\square$

Appendix C

Chapter 3: Proofs of Main Results

Proof of Lemma 1. Recall that the objective is to minimize $\mathcal{C}(\mathcal{T}_n)$ subject to $0 \leq t_1 \leq t_2 \leq \ldots \leq t_n \leq t_{n+1}$. We need to show that the optimal inspection policy cannot be a boundary solution. In Proposition 2 (Appendix A), we established that $t_n < t_{n+1} < \infty$. Suppose, by way of contradiction, that an inspection policy, $\mathcal{T}_n$, satisfies $t_m = t_{m-1}$ for some $1 \leq m \leq n$. We will show that this inspection policy cannot be optimal. Let $\mathcal{T}'_n = \{t'_j\}_{j=1}^{n+1}$ be a new inspection policy that performs an inspection at every time $\mathcal{T}_n$ does, but at time t_m, $\mathcal{T}'_n$ performs one less inspection than $\mathcal{T}_n$, and schedules this inspection at a later time. Using (3.3), it can be shown that for the new inspection policy, the probability of incurring the penalty does not increase while the expected discounted inspection cost decreases since this inspection is scheduled at a later time. Therefore, $\mathcal{T}'_n$ achieves lower expected discounted cost, from which we conclude that the optimal inspection policy is an interior solution. Thus, it should satisfy the first order necessary condition: $\frac{\partial \mathcal{L}_\alpha(\mathcal{T}_n)}{\partial t_j} = 0$ for each $1 \leq j \leq n+1$. By differentiating the expression of $\mathcal{L}_\alpha(\mathcal{T}_n)$ given in (3.3) with respect to t_j and setting it equal to zero, after simplification and rearrangement, we obtain for each $1 \leq j < n$:

$$\Big[\frac{F(t_j) - F(t_{j-1})}{f(t_j)} - \frac{1 - e^{-(\lambda+\theta)(t_{j+1}-t_j)}}{\lambda+\theta}\Big] - \frac{c_i}{k}\Big[\frac{1 - F(t_j)}{f(t_j)} + \frac{1}{\lambda+\theta}\Big] = 0, \tag{C.1}$$

where $k = \frac{\lambda}{\lambda+\theta}(c_p + r + \alpha) - (c_i + r + \alpha)$ by definition. Note that the left hand side of (A.4) is equal to $\Omega(t_j, x_j, x_{j+1})$. Similarly, we can show that the first order necessary conditions for $j = n$ and $j = n+1$ are equivalent to (3.7) and (3.8). □

Proof of Theorem 3. It suffices to show that if $x_{m+1}[t^*_{1,n}] > x_m[t^*_{1,n}]$ for some m, then $x_{m+2}[t^*_{1,n}] > x_{m+1}[t^*_{1,n}]$ follows. From (3.6) and (3.7), we have

$$\Omega(t_{m+1}[t^*_{1,n}], x_{m+1}[t^*_{1,n}], x_{m+2}[t^*_{1,n}]) \leq \Omega(t_m[t^*_{1,n}], x_m[t^*_{1,n}], x_{m+1}[t^*_{1,n}]) = 0, \tag{C.2}$$

where the inequality is strict if $m = n - 1$. Note that $t_{m+1}[t^*_{1,n}] > t_m[t^*_{1,n}]$ and $x_{m+1}[t^*_{1,n}] > x_m[t^*_{1,n}]$. Now, since $\Omega(t,x,y)$ is increasing in t, strictly increasing in x and strictly decreasing in y (Lemma 9 in Appendix A), we must have $x_{m+2}[t^*_{1,n}] > x_{m+1}[t^*_{1,n}]$ in order to satisfy (A.5). This completes the proof. □

Proof of Lemma 2. Since $x_1[t_{1,n}] = t_{1,n}$, from (3.6) for $j = 1$, we have $\Omega(t_{1,n}, t_{1,n}, x_2[t_{1,n}]) = 0$. To satisfy this equation, $x_2[t_{1,n}]$ must be strictly increasing in $t_{1,n}$ since $\Omega(t,x,y)$ is increasing in t, strictly increasing in x and strictly decreasing in y (Lemma 9 in Appendix A). Consequently, $t_2[t_{1,n}] = t_{1,n} + x_2[t_{1,n}]$ is strictly increasing in $t_{1,n}$. Now by inductive reasoning it follows that $t_j[t_1]$ and $x_j[t_1]$ are strictly increasing in t_1 for any $1 \leq j \leq n$. Now, from (3.7) for $j = n$, we have $\Omega(t_n[t_{1,n}], x_n[t_{1,n}], x_{n+1}[t_{1,n}]) = -\frac{c_i e^{-(\lambda+\theta)x_{n+1}}}{k(\lambda+\theta)} < 0$. Similar to the discussion above, since $\Omega(t,x,y)$ is increasing in t, strictly increasing in x and strictly decreasing in y, we conclude that $t_{n+1}[t_1]$ and $x_{n+1}[t_1]$ are also strictly increasing in t_1. □

Proof of Proposition 1. By Lemma 1, the optimal inspection policy must satisfy the first order necessary condition. Recall the necessary condition (3.8) is given by

$$\frac{F(t_n[t_{1,n}] + x_{n+1}[t_{1,n}]) - F(t_n[t_{1,n}])}{1 - F(t_n[t_{1,n}])} = \frac{\theta(r+\alpha)}{\lambda c_p}.$$

Now, by Lemma 2 and the fact that $\frac{F(t+x)-F(t)}{1-F(t)}$ is increasing in t and strictly increasing in x (by the increasing failure rate property of density f), the left hand side of this equality is strictly increasing in $t_{1,n}$. It follows that the value of $t_{1,n}$ satisfying the necessary condition is unique (which we denote by $t^*_{1,n}$). Now, using (3.6) and (3.7), the remaining inspection times ($t_j[t^*_{1,n}]$ for each $1 < j \leq n$) can be obtained recursively (and uniquely) by having $t^*_{1,n}$. Thus, we conclude the optimal inspection policy is unique. □

Appendix D

Chapter 3: Supporting Results

Proposition 2 *It is always optimal to schedule a replacement at a finite time t_{n+1} after the last inspection time, i.e., $t_n < t_{n+1} < +\infty$.*

Proof. Let $\widetilde{\mathcal{T}}_n$ be the inspection policy that does not schedule a replacement at a finite time (i.e., $t_{n+1} = +\infty$). First, we show that the best of such policies must schedule the inspection times $t_1, \ldots, t_n$ at finite times. We show that for any inspection policy that schedules m finite inspection times ($m < n$) of the allotted inspection opportunities, there exists a strictly better inspection policy that schedules $m+1$ inspection times. For any $1 \leq m < n$, let $\widetilde{\mathcal{T}}_m = \big\{t_j\big\}_{j=1}^{m}$ be an inspection policy with m inspection opportunities scheduled at finite times $t_1, t_2, \ldots, t_m$. Moreover, consider $\widetilde{\mathcal{T}}_{m+1}$ as an inspection policy with $m+1$ inspection opportunities where the first m inspection times $\{t_1, t_2, \ldots, t_m\}$ are the same as $\widetilde{\mathcal{T}}_m$, and we choose the $(m+1)^{th}$ inspection time t_{m+1} such that $1 > F(t_{m+1}) > F(t_m) + \frac{c_i}{k+c_i}\big(1 - F(t_m)\big)$. Such a value of t_{m+1} exists since $F(t_m) + \frac{c_i}{k+c_i}\big(1 - F(t_m)\big) < 1$ by the fact $k > 0$, and the $F(t_{m+1})$ is strictly increasing in t_{m+1} (since we assumed $f(t) > 0$) and approaches 1 when t_{m+1} is sufficiently large. Then, from (3.3) and after further simplification, we have

$$\mathcal{L}_\alpha(\widetilde{\mathcal{T}}_{m+1}) - \mathcal{L}_\alpha(\widetilde{\mathcal{T}}_m) = \\ \Big[-(k + c_i)\big(F(t_{m+1}) - F(t_m)\big) + c_i\big(1 - F(t_m)\big)\Big] e^{-(\lambda+\theta)t_{m+1}}.$$

By the choice of t_{m+1}, this expression is negative. Thus, the expected total cost can be (strictly) decreased by scheduling one more inspection at an appropriate time.

Now, for a general inspection policy $\mathcal{T}_n$, we can write after simplification:

$$\begin{aligned}\mathcal{L}_\alpha(\mathcal{T}_n) &= \mathcal{L}_\alpha(\widetilde{\mathcal{T}}_n) \\ &\quad + \int_{t_{n+1}}^{+\infty} \Big[-c_p\big(F(t) - F(t_n)\big) + \frac{\theta(r+\alpha)}{\lambda}\big(1 - F(t_n)\big)\Big] \lambda e^{-(\lambda+\theta)t} dt.\end{aligned} \tag{D.1}$$

Then, we choose a $t_{n+1} > t_n$ such that the integral in (D.1) is negative (note that for $t_{n+1} = +\infty$, the integral in (D.1) is equal to zero). It suffices to find a t_{n+1} that satisfies the following inequality:

$$F(t_{n+1}) > F(t_n) + \frac{\theta(r+\alpha)}{\lambda c_p}\big(1 - F(t_n)\big).$$

Such a value of t_{n+1} exists since $\frac{\theta(r+\alpha)}{\lambda c_p} = \frac{\theta(r+\alpha)}{\theta(r+\alpha)+(\lambda+\theta)(k+c_i)} < 1$ by the fact $k > 0$, and the $F(t_{n+1})$ is strictly increasing in t_{n+1} (since we assumed $f(t) > 0$) and approaches 1 when t_{n+1} is sufficiently large. Thus, an optimal inspection policy should schedule a replacement at a finite time after the last inspection. $\square$

Proposition 3 *For any n, an optimal inspection policy $\mathcal{T}_n^*$ exists.*

Proof. For some $\varepsilon > 0$, let t_{n+1}^M be a replacement time such that the integral in (D.1) is larger than $-\varepsilon$ for all $t_{n+1} > t_{n+1}^M$ and any value of t_n. Such t_{n+1}^M exists since the integral approaches 0 as $t_{n+1} \to \infty$.

By Proposition 2, the optimal inspection policy should schedule a replacement at a finite time. Now, consider an arbitrary inspection policy $\mathcal{T}_n$ such that $t_{n+1}^M < t_{n+1} < +\infty$ and let $\widetilde{\mathcal{T}}_n^*$ be the optimal inspection policy in the class of policies that do not schedule a replacement at a finite time and $t_n < t_{n+1}$. Such an optimal policy exists because $\mathcal{L}_\alpha(\widetilde{\mathcal{T}}_n)$ is continuous and attains its minimum over the following compact set:

$$\widetilde{D} = \Big\{0 \le t_j \le t_{n+1},\ t_j \ge t_{j-1} \text{ for all } 1 \le j \le n\Big\}.$$

Then, we obtain from (3.3) after simplification:

$$\begin{aligned}\mathcal{L}_\alpha(\mathcal{T}_n) &= \mathcal{L}_\alpha(\widetilde{\mathcal{T}}_n) \\ &\quad + \int_{t_{n+1}}^{+\infty} \Big[-c_p\big(F(t) - F(t_n)\big) + \frac{\theta(r+\alpha)}{\lambda}\big(1 - F(t_n)\big)\Big] \lambda e^{-(\lambda+\theta)t} dt \\ &\quad > \mathcal{L}_\alpha(\widetilde{\mathcal{T}}_n^*) - \varepsilon,\end{aligned}$$

where the first equality follows from (3.3), and the inequality comes from the choice of t_{n+1} and the fact that $\mathcal{T}_n^*$ minimizes $\mathcal{L}_\alpha(\widetilde{\mathcal{T}}_n)$. In other words, a policy that schedules inspections based on $\mathcal{T}_n^*$ and a replacement at t_{n+1}^M is better than any $\mathcal{T}_n$ with $t_{n+1} > t_{n+1}^M$. It implies that if the minimum of $\mathcal{L}_\alpha(\mathcal{T}_n)$ exists, it belongs to the following set:

$$D = \Big\{ 0 \leq t_j \leq t_{n+1}^M,\ t_j \geq t_{j-1} \ \text{ for all } 1 \leq j \leq n+1 \Big\}.$$

Since $\mathcal{L}_\alpha(\mathcal{T}_n)$ is continuous and D is compact, it follows that $\mathcal{L}_\alpha(\mathcal{T}_n)$ attains its minimum over the set D. Thus, an optimal inspection policy $\mathcal{T}_n^*$ exists. This completes the proof. □

Lemma 9 *For PF_2 densities f and g, and for $(t,x,y) \in \mathbb{R}_+^3$, the following statements hold:*

(a) $\Omega(t,x,y)$ is increasing in t.

(b) $\Omega(t,x,y)$ is strictly increasing in x for all y and for all t such that $t-x>0$.

(c) $\Omega(t,x,y)$ is strictly decreasing in y for all t and x.

Proof. We only prove parts (a) and (b). Proof of part (c) is similar to part (b).
(a) Since the first square bracket of (3.5) is increasing in t (Theorem 3 in Barlow et al. [11]), the second square bracket of (3.5) is decreasing in t (Corollary 3.1 in Barlow et al. [11]), and $k>0$, it follows that $\Omega(t,x,y)$ is increasing in t.
(b) Since the density function f satisfies $f(a)>0$ for any $a>0$, we have $\frac{\partial \Omega(t,x,y)}{\partial x} = \frac{f(t-x)}{f(t)} > 0$. Thus, $\Omega(t,x,y)$ is strictly increasing in x. □

Appendix E

Chapter 3: Solution Algorithm

We first present the approximation algorithm and then we explain its logic. The algorithm starts with an initial search interval $(\underline{t}, \overline{t})$, for which we can conveniently choose $\underline{t} = 0$ and $\overline{t}$ a sufficiently large number (such that the probability that the failure occurs before $\overline{t}$ is close to 1). Then, the following steps are followed (note for the sake of clarity, we do not show the dependence of the inspection times t_j on the first inspection time):

1. Given the search interval $(\underline{t}, \overline{t})$, set $t_1 = \frac{\underline{t}+\overline{t}}{2}$, $x_1 = t_1$, and $j = 1$.
2. If $\overline{t} - \underline{t} < 2\epsilon$, STOP and return $t_{1,n}^{out} = \frac{\overline{t}+\underline{t}}{2}$ as the optimal first inspection time $t_{1,n}^*$.
3. While $j < n$, take the following steps:
   (a) Compute $Q := \frac{F(t_j)-F(t_j-x_j)}{f(t_j)} - \frac{c_i}{k}\left[\frac{1-F(t_j)}{f(t_j)} + \frac{1}{\lambda+\theta}\right]$.
   (b) If $Q \leq 0$, then set $\underline{t} = t_{1,n}$, and go back to Step 1.
   (c) If $Q \geq \frac{1}{\lambda+\theta}$, then set $\overline{t} = t_{1,n}$, and go back to Step 1.
   (d) Find x_{j+1} that satisfies $\frac{1-e^{-(\lambda+\theta)x_{j+1}}}{\lambda+\theta} = Q$. Increase j by 1 and go back to Step 3.
4. Compute $Q := \frac{F(t_n)-F(t_n-x_n)}{f(t_n)} - \frac{c_i}{k}\left[\frac{1-F(t_n)}{f(t_n)} + \frac{1}{\lambda+\theta}\right]$.
   (a) If $Q \leq -\frac{1}{\lambda+\theta}\frac{c_i}{k}$, then set $\underline{t} = t_{1,n}$, and go back to Step 1.
   (b) If $Q \geq \frac{1}{\lambda+\theta}$, then set $\overline{t} = t_{1,n}$, and go back to Step 1.
   (c) Find x_{n+1} that satisfies $\frac{1-(1+\frac{c_i}{k})e^{-(\lambda+\theta)x_{n+1}}}{\lambda+\theta} = Q$, and go to Step 5.
5. Compute $Q := \frac{F(t_n+x_{n+1})-F(t_n)}{1-F(t_n)}$.

(a) If $Q < \frac{\theta(r+\alpha)}{\lambda c_p}$, then set $\underline{t} = t_{1,n}$, and go back to Step 1.

(b) If $Q > \frac{\theta(r+\alpha)}{\lambda c_p}$, then set $\bar{t} = t_{1,n}$, and go back to Step 1.

In each iteration of the algorithm, we choose the midpoint of the search interval as the first inspection time t_1. Then, using (3.6) and (3.7), we try to recursively find the other inspection times $\{t_2, \ldots, t_n\}$ by solving the equation $\frac{1-e^{-(\lambda+\theta)x_{j+1}}}{\lambda+\theta} = Q$ for each $j = 1, \ldots, n-1$, which is equivalent to (3.6) (see Step 3(d)). Now, since the choice of t_1 is not necessarily optimal, for some j, we might not be able to find any x_{j+1} that satisfies the above equation. Since Q is increasing in t_j and strictly increasing in x_j (see the proof of Lemma 9(a)-(b)), it follows that Q is strictly increasing in t_1 (by Lemma 2). Therefore, by checking whether the value of Q is too large or too small, we can decide whether to update either the lower bound or the upper bound of the search interval (see Steps 3(b) and 3(c)). The same logic applies to the case $j = n$ and $j = n + 1$ (see Steps 4 and 5). The length of the next iteration's search interval is half of the previous iteration's search interval. We continue in this manner until we conclude that $t^*_{1,n}$ belongs to a search interval of a required length 2ϵ. Then, the approximation algorithm returns the midpoint of the last search interval as $t^*_{1,n}$.