ABSTRACT

ZHANG, YUAN. A Partially Observable Markov Decision Process for Optimal Design of Surveillance Policies for Bladder Cancer. (Under the direction of Dr. Brian Denton.)

Bladder Cancer is the fourth most common cancer in men and eighth in women in the United States. For patients with a history of bladder cancer, the probability of recurrence at one year ranges from 15% to 70%; and the probability of progression to high risk muscle invasive bladder cancer at 5 years ranges from 7% to 40%, depending on the patient's particular risk factors. Cystoscopy is regarded as the gold standard for surveillance of bladder cancer recurrence and progression. However, no consensus exists about the best frequency of follow-up cystoscopy for patients with a history of low grade Ta disease. In this thesis we use stochastic models to investigate policies for bladder cancer surveillance. First, we formulate a partially observable Markov model. The model includes states defining stages of bladder cancer, the effects of treatment, death from bladder cancer, and all other cause mortality. Simulation is used to compare published recommendations for bladder cancer surveillance policies based on expected quality adjusted life years (QALYs) over the patient's lifetime. We compare the American Urological Association (AUA) guideline, the European Association of Urology (EAU) guideline, and several other policies. Next, we extend our model to a partially observable Markov decision process (POMDP) to determine the optimal surveillance policy. We present a series of computational experiments. Results show that age and comorbidity significantly affect the optimal surveillance policy. We find that younger patients should have more intensive surveillance than older patients and patients having comorbidity should have less intensive surveillance. We perform sensitivity analysis to evaluate the influence of model input parameters. Among them we find that disutility of cystoscopy has a significant influence on the optimal surveillance policy. In general, the lower the disutility of cystoscopy, the more intensively surveillance should be performed. Finally, we extend our initial POMDP model to incorporate a new urine based biomarker test into the surveillance process. We study the incremental benefit of the optimal policy that includes a biomarker test and cystoscopy over the optimal policy based on cystoscopy alone. We also compare the optimal policy to easy-to-implement heuristic policies using a biomarker to direct the frequency of cystoscopies. We find that introduction of a biomarker does not significantly improve the optimal policy based on cystoscopy alone; however, biomarkers may significantly improve the heuristic policies we investigated.

A Partially Observable Markov Decision Process for Optimal Design of Surveillance Policies for Bladder Cancer

by
Yuan Zhang

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Operations Research

Raleigh, North Carolina

2012

APPROVED BY:

Dr. Michael Devetsikiotis	Dr. Thom Hodgson
Dr. Julie Ivy	Dr. Brian Denton Chair of Advisory Committee

DEDICATION

To my mother and late father.

BIOGRAPHY

The author, Yuan Zhang, was born in a beautiful town located beside the Jialing river in Sichuan province of China on June 6, 1983. Several months later, he moved to Baoding, Hebei province and had lived there till he was 18. He studied in the department of Mathematics Zhejiang Univeristy from 2001 to 2005.

He came to North Carolina State University for his PhD study in the graduate program of Operations Research, in 2006. Shortly after his dissertation defence, Yuan joined the Department of Quantitative Health Sciences at the University of Massachusetts Medical School as an instructor in January 2012. When he is not working, he enjoys travelling, running, swimming and skiing.

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to my advisor Professor Brian Denton for continuous support of my Ph.D research. His motivation, patience, enthusiasm, and immense knowledge has embarked me on an academic career in operations research. I could not have imagined having a better advisor and mentor for my Ph.D study.

Besides my advisor, I would like to express my gratitude to Dr. Matthew Nielsen, a urologist at University of North Carolina, Chapel Hill. Dr. Nielsen gave me tremendous help during my Ph.D research, including guidance on reviewing the medical literature, building quantitative models, and validating the results.

My sincere thanks also goes to Professor Michael Devetsikiotis and Dr. Andy Rindos for offering me the research opportunity at Professor Devetsikiotis' network research performance group and Dr. Rindos' for the opportunity to work at the Center for Advanced Studies at IBM.

I would also like to thank the rest of my thesis committee: Prof. Julie Ivy and Prof. Thom J. Hodgson for their encouragement and insightful comments during the course of my PhD study.

I thank my fellow labmates at NCSU for making my research less painful and my life more enjoyable. My special thanks goes to Daniel Underwood for teaching me LINUX, and miscellaneous programming skills. I also thank Daniel, Akshay V.S., Jennifer Mason, Bjorn Berg , Jingyu Zhang and Yuanhui Zhang for their enlighten discussion on my research problems and proofreading my dissertation manuscript.

Last but not least, I would like to thank my family for their love, support and encouragement for me to achieve this goal.

This thesis is based in part upon work supported by the National Science Foundation under Grant Number CMMI 0844511.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1

Introduction

Bladder cancer is the fourth most common cancer in men and the eleventh most common in women, accounting for 3.2% of all men who died of cancer between 2001 and 2005 [31]. The U.S. National Cancer Institutes (NCI) estimated that there were 70,530 new cases of bladder cancer (52,760 men and 17,770 women) diagnosed and 14,680 deaths caused by bladder cancer (10,410 men and 4,270 women) in 2010 [34].

Bladder cancer occurs in two clinically significant forms: (1) non-muscle invasive bladder cancer (NMIBC) and (2) muscle invasive bladder cancer (MIBC). MIBC typically results in worse outcomes than NMIBC, often resulting in death from bladder cancer. NMIBC represents a heterogeneous group of tumors with completely different ontological outcomes. Low grade NMIBC has a modest recurrence rate and a very low risk for progression. High grade NMIBC, on the other hand, is associated with much higher recurrence, progression, and mortality rates than low grade NMIBC [14]. According to the World Health Organization (WHO) cancer report (2008) [7], approximately 75% of individuals with bladder cancer have NMIBC. Among patients with NMIBC approximately 70% of tumours are low grade [38]. A significant proportion of patients with NMIBC (with varying estimates from 31-78%) have at least one recurrence during a 5-year period [66].

Definitive treatment such as cystectomy (surgical removal of the bladder or part of the bladder) is often recommended for patients with MIBC or high grade bladder cancer. For patients with low grade NMIBC, on the other hand, regular surveillance is recommended. Cystoscopy is currently the gold standard for bladder cancer surveillance. However, the procedure can be painful, and a source of anxiety for patients [5] [29]. Current guidelines are not consistent in their recommendations about the frequency of cystoscopy for low risk bladder cancer patients [4] [54] [52] [24]. The International Bladder Cancer Group (IBCG) compared guideline recommendations of the European Association of Urology (EAU, 2009) [4], the First International Consultation on Bladder Tumors (FICBT, 2005) [54], the National Comprehensive Cancer Net-

work (NCCN, 2010) [52], and the American Urological Association (AUA, 2007) [24]. The IBCG concluded that there is currently no consensus on optimal surveillance for low risk patients [55]. Nevertheless, a study by Schrag et al. [60] showed that many urologists and patients seem to have chosen a lower cystoscopy frequency than those guidelines.

This thesis investigates the design of surveillance policies for low risk bladder cancer patients. First, we use Monte-Carlo simulation to compare international guidelines for low risk bladder cancer surveillance. Next, we extend our model to a partially observable Markov decision process (POMDP) model to investigate the optimal surveillance policies. Finally, we extend the POMDP model to incorporate new urine based biomarker tests into the surveillance process. We analyze the incremental benefit of biomarker tests for improving the optimal surveillance policies and we investigate easy-to-implement heuristic surveillance schedules that are near to optimal. Following is a detailed description of each of the chapters of this thesis.

Chapter 2 provides some background on bladder cancer and motivation for studying low risk bladder cancer surveillance. We describe some important statistics related to bladder cancer. We discuss risk factors, methods for bladder cancer detection, and common forms of treatment. We also provide a methodological review of the POMDP literature, including theoretical properties of POMDPs and algorithms for solving POMDPs. Finally we review some recent applications of POMDPs to medical decision making.

In Chapter 3 we describe a partially observable Markov model based on states that define patient risk levels associated with recurrence and progression of bladder cancer. The model includes states defining the effects of treatment, death from bladder cancer, and all other cause mortality. The model is partially observable in that the precise health state of the patient is unknown in the absence of cystoscopy test results. International bladder cancer surveillance guidelines are compared with alternative surveillance policies based on expected quality adjusted life years (QALYs) over the patient's lifetime. Monte-Carlo sampling is used to generate 100,000 sample paths for each surveillance policy to estimate expected QALYs. Sensitivity analysis is performed on the patient's disutility associated with cystoscopy, bladder cancer mortality, and all other cause mortality. Computational results show that age and comorbidity affect the ranking of the best surveillance policies. We find that younger patients should have more intensive surveillance than older patients. Patients having comorbidity should have less intensive surveillance. Model parameters also affect the optimal strategy. The lower the disutility of cystoscopy on patients, the more intensively surveillance should be performed.

In Chapter 4 we formulate a POMDP model based on states that define patient risk levels associated with recurrence and progression of bladder cancer to study the optimal surveillance policy that maximizes expected QALYs. Optimal policies are computed using an exact method called *incremental pruning*. We compare the optimal policy to the current international guidelines studied in Chapter 3 for male and female patients respectively. We find the optimal

policy for typical male patients can result in an expected gain of 0.4 QALYs over the EAU and AUA guidelines for our base case computational experiments. Sensitivity analysis is performed on the patient's disutility associated with cystoscopy, bladder cancer mortality, and all other cause mortality. Results show that age, gender, and comorbidity significantly affect the optimal surveillance policy. We find that younger patients should have more intensive surveillance than older patients. Patients having comorbidity should have less intensive surveillance. We perform sensitivity analysis to evaluate the influence of model input parameters. Among them we find that disutility of cystoscopy has a significant influence on the optimal surveillance policy. In general, the lower the disutility of cystoscopy, the more intensively surveillance should be performed.

In Chapter 5 we extend the POMDP model of Chapter 4 by incorporating a urine based biomarker test. In addition to the standard surveillance protocol that uses cystoscopy alone, we consider an alternative surveillance protocol in which a biomarker test is performed first, with a positive result triggering a follow-up cystoscopy. We investigate the optimal surveillance policy using both the standard protocol and the alternative protocol by maximizing the total expected QALYs. We analyze the incremental benefit of using a biomarker to direct the frequency of diagnostic tests by comparing the outcomes resulting from an easy-to-implement heuristic policy with and without using the biomarker as part of the surveillance protocol.

In Chapter 6 we conclude with a summary of the most significant findings from Chapter 3, 4, and 5. We discuss some of the limitations of our work, and opportunities for future research related to bladder cancer surveillance.

Chapter 2

Background and Literature Review

2.1 Introduction

In this chapter we first provide some background on bladder cancer and motivation for studying low risk bladder cancer surveillance. In Section 2.3 we review the standard formulation of a POMDP. In Section 2.4 we provide a literature review of theoretical methods and algorithms for solving POMDPs. In Section 2.5 we provide a review of applications of POMDP to medical decision making. Finally, in Section 2.6 we summarize the contributions of this thesis to the literature.

2.2 Bladder Cancer Background

Early detection of bladder cancer can reduce disease related mortality. According to the National Cancer Institute's Surveillance, Epidemiology and End Results (SEER) program [31], 50% of urinary bladder cancer cases are diagnosed while the cancer is only in the layer of cells in which it began (in situ stage); 36% are diagnosed while the cancer is still confined to the primary site (localized stage). These early stage bladder tumors are also called NMIBCs, as they have not invaded the muscle layer of the bladder. For NMIBCs effective medical and surgical treatment options are available which can reduce or eliminate the possibilities of progression to later stages.

The severity of bladder cancer is characterized by the degree to which it has spread in the bladder and the body. NMIBC is limited to the innermost linings of the bladder. Most patients with NMIBC have very low risk of progression to MIBC, but the 5-year relative survival rate drops from 95% to 36.2% when the cancer has spread to regional lymph nodes or beyond the primary site; even worse it drops to 5.8% when the cancer has metastasized. However, NMIBC represents a heterogeneous group of tumors with different ontological outcomes. Low grade

NMIBC has a modest recurrence rate but a very low risk for progression. High grade NMIBC, on the other hand, is associated with significantly higher recurrence, progression, and mortality rates [14]. In spite of these differences current recommendations for NMIBC are "one size fits all".

2.2.1 Bladder Cancer Detection

Most cancers are detected when patients present with symptoms of the disease. In the case of bladder cancer, the most common cancer symptom is haematuria, the presence of red blood cells in the urine. It is the most common finding in Ta and T1 NMIBCs because they do not cause bladder pain and rarely present with bladder irritation, dysuria or urgency [4]. Once a patient has been diagnosed and treated with bladder cancer, cystoscopy is used to perform regular surveillance. Cystoscopy is a diagnostic test in which a urologist looks inside of the bladder and urethra with a thin lighted tube called a cystoscope. A description of the cystoscopy findings include the site, size, number, and appearance (papillary or solid) of the tumors as well as a description of mucosal abnormalities [4]. Urine cytology is also commonly used in combination with cystoscopy. It involves the examination of a urine specimen for exfoliated cancer cells. However it is of limited value for low grade tumors due to operator dependency and low sensitivity. Therefore it is normally combined with cystoscopy.

The limitations of cytology and the invasiveness of cystoscopy for detecting bladder cancer have generated interest in the development of new urine based biomarker tests. These simple tests involve the use of urine samples to detect the recurrence of certain types of bladder tumors. According to van Rhijn, et al. (2005) [67], microsatellite analysis, CYFRA21-1 and LewisX are the most promising non-FDA approved urine based biomarkers while NMP22, ImmunoCyt and FISH are the best FDA approved tests for surveillance.

NMP22 is a point-of-care biomarker which does not require expert analysis or laboratory time. The cost of NMP22 is less than half that of cytology [23]. NMP22 has much higher sensitivity than cytology, but its specificity is lower [23] [53]. Studies show the combination of NMP22 and cystoscopy can identify 99% of all malignancies versus 91.3% with cystoscopy alone [23] [68]. According to Grossman, et al [23], Immunocyt and FISH have some limitations compared with NMP22. Immunocyt is FDA-approved for surveillance only in conjunction with traditional urine cytology; FISH is the only test other than the NMP22 biomarker that is FDA-approved for use in diagnosis (in patients with hematuria only) as well as surveillance, but published sensitivity and specificity were not calculated from the target population. Unlike the NMP22 test, costs for Immunocyt and FISH are equivalent to or higher than urine cytology, thereby increasing the cost of cancer detection.

Fritsche, el al [20] summarized several potential applications of urine based biomarker tests

in patient surveillance, including serial testing to detect recurrent disease and as an adjunct to urine cytology to direct the frequency of cystoscopy evaluation in the follow-up of patients with bladder cancer. There are many possible ways to combine urine based biomarkers with cystoscopy and/or cytology in bladder cancer surveillance. Mowatt, et al [53] proposed two options based on advice from clinical experts. Option one is to use one test tool (flexible cystoscopy, cytology or biomarker) as an initial test; a positive result from the initial test would qualify the patient for cystoscopy. Option two is to use two tests (either flexible cystoscopy and a biomarker, or flexible cystoscopy and cytology) initially, and a positive result of both the initial tests would qualify the patient for cystoscopy.

Even though most proposed biomarkers have higher sensitivities than cytology, none of them has comparable sensitivity with cystoscopy; therefore at the present time it seems unlikely cystoscopy could be entirely replaced by biomarkers. However some studies have suggested that biomarker tests could be used to help determine the frequency of cystoscopy evaluation in the follow-up of patients with bladder cancer [20]. For example, an alternative surveillance protocol could alternate cystoscopies with biomarker tests. This could reduce the number of cystoscopies over a patient's lifetime but possibly increase the risk of progression since the lower sensitivity of the urine based biomarker test may result in failure to detect tumors. Based on a systematic review, van Rhijn et al.(2005) [67] concluded that current evidence is insufficient to determine if or how to use a urine based biomarker test for bladder cancer surveillance.

2.2.2 Bladder Cancer Treatment

Once diagnosed, NMIBCs can often be removed by surgery, called transurethral resection of bladder tumors (TURBT). This involves passing a high-frequency electric current through a wire inserted through the cystoscope, removing and burning cancer cells on the bladder wall. Patients with high risk bladder tumors are often treated with TURBT combined with chemotherapy, which may be delivered by mouth, intravenously, or instilled directly into the bladder. Chemotherapy may cause many side effects including bloody urine and bladder irritation causing increased urination frequency, urgency, pain and/or burning with urination. Bacillus Calmette-Guerin solution (BCG) is another way to treat bladder cancer. It contains weakened bacteria that stimulate the immune system to kill cancer cells in the bladder. The physician uses a catheter (a thin, flexible tube) to put the solution in the bladder, and the patient typically hold the solution in the bladder for about two hours. This treatment is usually done once a week for at least six weeks. Side effects may include irritation of the bladder, urination urgency, and urination frequency.

MIBC has a high risk of progression to metastatic cancer. As a result, the standard surgical treatment for MIBC is cystectomy, surgical removal of all or part of the urinary bladder.

Approximately 50% of NMIBC patients will need cystectomy eventually, and one-third are at risk over a 15-20 year period of dying of bladder cancer [27].

2.2.3 Bladder Cancer Risk

For patients that have been diagnosed and treated for early stage bladder cancer, the risk of recurrence and progression of the disease is an important consideration in designing a surveillance policy. Prognostic factors for recurrence and progression include the number of tumors, tumor size, prior recurrence rate, stage, and grade. The EAU [66] developed a simple scoring system based on these prognostic factors, as described in Table 2.1. A recurrence score can be calculated for each patient, from 0 (best prognosis) to 17 (worst prognosis). Patients can be then divided into four groups according to their score, as shown in Table 2.2. For example, a patient with a primary small ($\leq$ 3 cm) tumor of the lowest stage and the lowest grade would have a recurrence score of 0 and a progression score of 0, which indicate a recurrence rate of 15% and a progression rate of 0.2%, at one year, respectively. If the patient has a recurrent small ($\leq$ 3cm) tumor of the same stage and grade after surgical removal of the primary tumor, his 1 year recurrence rate and 1 year progression rate would increase to 24% and 1% respectively, indicated by the new recurrence score of 2 and the new progression score of 2.

2.3 Partially Observable Markov Decision Processes

A Markov decision process (MDP) is a sequential decision process in which an underlying stochastic process, a Markov process, is combined with a set of feasible actions at each state, and rewards that depend on the action and state. MDPs have been applied to several types of medical decision making problems. POMDPs are well known extensions of (completely observable) MDPs to the case in which some states are not directly observable. POMDPs assume the state of a specific system is not known with certainty, but can be described probabilistically by a *belief state*. Medical diagnosis fits this context well since it often prone to errors due to imperfect sensitivity and specificity of diagnostic tests. Since the exact health state of the patient never completely reveals itself, partial observability is an integral part of medical decision making.

In this section, we describe the standard formulation of a finite state POMDP. The following description uses notation similar to that of Monahan (1982) [51]. Let X_t define a *core state* of the unobservable process. It is a random variable which takes on values in the finite set $S \equiv \{1, ..., N\}$. The stochastic process $\{X_t, t \in T\}$, called the *core process*, is assumed to be a finite state Markov chain. Associated with X_t is a random variable Y_t that defines the *observation state*, which takes on values in the finite observation set $\Theta \equiv \{1, ..., M\}$. The stochastic process $\{Y_t, t \in T\}$ is called the *observation process*.

Table 2.1: The scoring system proposed by the EAU to calculate the recurrence and progression scores for bladder cancer [66].

Factor	**Recurrence**	**Progression**
Number of tumors		
Single	0	0
2 to 7	3	3
≥ 8	6	3
Tumor size		
< 3 cm	0	0
≥ 3 cm	3	3
Prior recurrence rate		
Primary	0	0
≤ 1 rec/yr	2	2
> 1 rec/yr	4	2
Stage		
Ta	0	0
T1	1	4
Ta and CIS	1	6
T1 and CIS	2	10
Grade		
G1	0	0
G2	1	0
G3	2	5
Total score	0-17	0-23

Table 2.2: Probability of recurrence and progression according to recurrence and progression scores [66].

Recurrence score	**Probability of recurrence at 1 year (95% CI)**	**Probability of recurrence at 5 years (95% CI)**
0	15% (10%, 19%)	31% (24%, 37%)
1 to 4	24% (21%, 26%)	46% (42%, 49%)
5 to 9	38% (35%, 41%)	62% (58%, 65%)
10 to 17	61% (55%, 67%)	78% (73%, 84%)
Progression score	**Probability of progression at 1 year (95% CI)**	**Probability of progression at 5 years (95% CI)**
0	0.2% (0%, 0.7%)	0.8% (0%, 1.7%)
2 to 6	1.0% (.4%, 1.6%)	6% (5%, 8%)
7 to 13	5% (4%, 7%)	17% (14%, 20%)
14 to 23	17% (10%, 24%)	45% (35%, 55%)

Let a_t be a decision variable, called an *action*, that takes on values in a finite set A_t. Assume that the decision maker can control the observation process by choosing actions $a_t \in A_t$. Let $q_t(\theta_t|s_t, a_t)$ denote the probability of $Y_t = \theta_t$ conditioning on $X_t = s_t$ given choosing action a_t at epoch t. $Q_t(a_t)$ is called the *information matrix* conditioned on action a_t, with elements $q_t(\theta_t|s_t, a_t)$. Assume the core process is related to the observation. Let $p_t(s_{t+1}|s_t, \theta_t)$ denote the core state transition probability from core state s_t to s_{t+1} at time epoch t given observation θ_t, and the matrix $P_t(\theta_t)$ is the transition probability matrix conditioning on observation θ_t, with elements $p_t(s_{t+1}|s_t, \theta_t)$.

At each epoch t, the information available for decision making is denoted by η_t. Define $\pi_t(s_t) \equiv Pr\{X_t = s_t \mid \eta_t\}$ to be the probability that $X_t = s_t$ given the available information η_t, and the vector $\pi_t = (\pi_t(1), ..., \pi_t(N))$ is called the *belief state*.

We assign a real number $r_t(s_t, a_t)$, called an *immediate reward*, to each core state s_t after taking action a_t at time epoch t. The vector $r_t(a_t) = (r_t(1, a_t), ..., r_t(N, a_t))^T$ is called the immediate reward vector. Thus, the expected immediate reward at a belief state π_t after taking action a_t can be written as $r_t(\pi_t, a_t) = \sum_{s_t \in S} \pi_t(s_t) r_t(s_t, a_t) = \pi_t r_t(a_t)$.

Without loss of generality, we assume that the rewards are discounted annually by $\lambda \in [0, 1]$. At each epoch, the expected immediate reward and the expected discounted future rewards are used to determine the optimal decision. The optimality equations at each epoch can be written as:

$$v_t(\pi_t) = \max_{a_t \in A_t} \left\{ r_t(\pi_t, a_t) + \lambda \sum_{\theta_t \in \Theta} v_{t+1}(\pi_{t+1}) \overline{p}_t(\theta_t|\pi_t, a_t) \right\}, \forall (t, \pi_t), t = 1, ..., T-1 \tag{2.1}$$

and

$$v_T(\pi_T) = \sum_{s_T \in S} \pi_T(s_T) r_T(s_T), \forall \pi_T. \tag{2.2}$$

The optimal action in epoch t at belief state π_t can be written as:

$$a_t^*(\pi_t) = \arg\max_{a_t \in A_t} \left\{ r_t(\pi_t, a_t) + \lambda \sum_{\theta_t \in \Theta} v_{t+1}(\pi_{t+1}) \overline{p}_t(\theta_t|\pi_t, a_t) \right\}, \tag{2.3}$$

where

$$\overline{p}_t(\theta_t|\pi_t, a_t) = \sum_{s_t \in S} \pi_t(s_t) q_t(\theta_t|s_t, a_t) \tag{2.4}$$

denotes the probability of observing θ_t given action a_t is taken at belief state π_t at epoch t. In

equation 2.1 and equation 2.3, π_{t+1} is determined by π_t, a_t, and θ_t by using Bayesian updates defined by the following formula:

$$\pi_{t+1}(s_{t+1}) = \frac{\sum_{s_t \in S} \pi_t(s_t) q_t(\theta_t|s_t, a_t) p_t(s_{t+1}|s_t, \theta_t)}{\sum_{s_t \in S} \pi_t(s_t) q_t(\theta_t|s_t, a_t)} \tag{2.5}$$

where $\pi_{t+1}(s_{t+1})$, the component of the belief vector, π_{t+1}, is a function of θ_t, a_t, and π_t. Thus equation 2.5 provides a means to update the belief state of the core process based on the prior belief state and the most recent action and most recent observation.

2.4 POMDP Methods

A number of exact methods and approximate solution methods for POMDPs have been proposed over the last 40 years. In this subsection we provide a brief review of exact methods directly related to this thesis. More general reviews can be found in [51], [45], [69], [35], [10], and [56].

Sondik [63] proved a number of properties of POMDPs which are central to solution methods. He showed that the optimal value function $v_t(\cdot)$, as defined in equation 2.1, is piecewise linear and convex. Therefore it can be expressed as follows: $v_t(\pi_t) = \max_{\alpha_t \in \Omega_t} \{\pi_t \alpha_t\}$, where Ω_t is a finite set of n-dimensional vectors, called the *α-vector set* that construct $v_t(\cdot)$. At epoch t, Ω_t can be recursively determined by substituting $v_{t+1}(\pi_{t+1}) = \max_{\alpha_{t+1} \in \Omega_{t+1}} \{\pi_{t+1} \alpha_{t+1}\}$ into equation 2.1, which results in the following:

$$v_t(\pi_t) = \max_{a_t \in A_t} \left\{ \pi_t r_t(a_t) + \lambda \sum_{\theta_t \in \Theta} \max_{\alpha_{t+1} \in \Omega_{t+1}} \{\pi_{t+1} \alpha_{t+1}\} \overline{p}_t(\theta_t|\pi_t, a_t) \right\}. \tag{2.6}$$

By plugging in equations 2.5 and 2.4, equation 2.6 can be rewritten as:

$$v_t(\pi_t) = \max_{a_t \in A_t} \left\{ \pi_t r_t(a_t) + \lambda \sum_{\theta_t \in \Theta} \max_{\alpha_{t+1} \in \Omega_{t+1}} \{\pi_t Q_t(\theta_t, a_t) P_t(\theta_t) \alpha_{t+1}\} \right\}, \tag{2.7}$$

where $Q_t(\theta_t, a_t)$ is defined as the following $N \times N$ diagonal matrix:

$$Q_t(\theta_t, a_t) = \begin{bmatrix} q_t(\theta_t|1, a_t) & & 0 \\ & \ddots & \\ 0 & & q_t(\theta_t|N, a_t) \end{bmatrix}.$$

Equation 2.7 can be rewritten in terms of α-vectors as:

$$v_t(\pi_t) = \max_{a_t \in A_t} \left\{ \pi_t \{ r_t(a_t) + \sum_{\theta_t \in \Theta} \max_{\alpha_t \in \Omega_{t+1}} \lambda Q_t(\theta_t, a_t) P_t(\theta_t) \alpha_{t+1} \} \right\}. \tag{2.8}$$

Given the above properties, the problem of solving a POMDP is equivalent to finding the α-vector set that describes $v_t(\cdot)$. Sondik proposed the first exact algorithm for solving a finite horizon POMDP in 1971 [63]. His algorithm, called *one-pass*, constructs Ω_t from the previous α-vector set Ω_{t+1} as follows:

$$\Omega_t = \left\{ \alpha_t = r_t(a_t) + \lambda \sum_{\theta_t \in \Theta} Q_t(\theta_t, a_t) P_t(\theta_t) \alpha_{t+1}^{\theta_t} \mid \forall a_t \in A_t, \forall \alpha_{t+1}^{\theta_t} \in \Omega_{t+1} \right\}.$$

Sondik's one-pass algorithm is often unable to solve POMDPs because of the exponential growth in the number of α-vectors in Ω_t as $|\Omega_t| = |A_t||\Omega_{t+1}|^M$. However, not all α-vectors in Ω_t are useful when determining the optimal value function, $v_t(\cdot)$, as some are dominated by other vectors in Ω_t. Thus, Monahan (1982) [51] proposed to prune Ω_t to its minimal representation Ω_t^* by removing dominated α-vectors. An α-vector $\alpha \in \Omega_t$ can be evaluated to determine whether or not is dominated by solving the following linear program $LP(\alpha, \Omega_t)$:

$$\begin{aligned} \text{Min} \quad & z(\pi, y) = y - \sum_{s \in S} \pi(s) \alpha(s) \\ \text{subject to} \quad & \sum_{s \in S} \pi(s) \alpha_t(s) \leq y, \forall \alpha_t \in \Omega_t \\ & \sum_{s \in S} \pi(s) = 1 \\ & 1 \geq \pi(s) \geq 0, \forall s \in S. \end{aligned} \tag{2.9}$$

If the optimal objective value $z(\pi^*, y^*) > 0$, then α is dominated by some vectors in Ω_t. Otherwise if $z(\pi^*, y^*) = 0$, then α supports the value function, $v_t(\cdot)$, at belief point π^*. This pruning process affects the running time by reducing the computation complexity from $|A_t||\Omega_{t+1}|^M$ to

$|A_t||\Omega^*_{t+1}|^M$, where Ω^*_{t+1} denotes the minimal representation of Ω_{t+1}. However this is done at the expense of solving $|A_t||\Omega^*_{t+1}|^M$ LPs at each epoch t. It has been shown that Monahan's algorithm remains exponential in the worst case and intractable for large size problems [10].

Rather than generating a large α-vector set and then pruning it to its minimal representation, Cheng (1988) [12] proposed an algorithm, called *linear support*, to compute $v_t(\cdot)$ by generating only the support α-vectors of $v_t(\cdot)$. Given a belief state π, it is trivial to see that the α-vector

$$\alpha^*_t(\pi) = \arg\max_{\alpha_t} \left\{ \pi\alpha_t \mid \alpha_t = r_t(a_t) + \lambda \sum_{\theta_t \in \Theta} Q_t(\theta_t, a_t) P_t(\theta_t) \alpha^{\theta_t}_{t+1}, \forall a_t \in A_t, \forall \alpha^{\theta_t}_{t+1} \in \Omega_{t+1} \right\}$$

is a support vector of $v_t(\cdot)$ at belief π, where π is called a *witness point* for $\alpha^*_t(\pi)$. Therefore, finding support α-vectors is equivalent to finding a set of corresponding witness points. Cheng proved that given a subset of support α-vectors of $v_t(\cdot)$, Ω'_t, one can search for witnesses of the missing support α-vectors among the vertices of the convex hull constructed by Ω'_t. A vertex π is a witness point for some missing support α-vector if $\alpha^*_t(\pi) \notin \Omega'_t$. Therefore, Ω'_t is incrementally augmented by finding a witness π and adding $\alpha^*_t(\pi)$ to Ω_t. Cheng's linear support algorithm is exponential in the worst case as the convex hull of a given Ω_t may have an exponential number of vertices [44].

Rather than constructing a support α-vector set to represent $v_t(\cdot)$, Kaelbling, Littman and Cassandra (1994) [43] proposed an algorithm, called *the witness algorithm*, to solve $v_t(\cdot)$ by concentrating on representing the conditional value function $v_t(\cdot \mid a_t)$ for each action $a_t \in A_t$ at a time, where

$$v_t(\pi_t \mid a_t) = \left\{ \pi_t r_t(a_t) + \lambda \sum_{\theta_t \in \Theta} \max_{\alpha_{t+1} \in \Omega_{t+1}} \{\pi_{t+1}\alpha_{t+1}\} \overline{p}_t(\theta_t | \pi_t, a_t) \right\},$$

and then combining $v_t(\cdot \mid a_t)$ to write the value function as $v_t(\cdot) = \sum_{a_t \in A_t} v_t(\cdot \mid a_t)$. Like Cheng's linear support algorithm, the witness algorithm starts with a subset of support α-vectors of $v_t(\cdot \mid a_t)$, Ω_{a_t}, and incrementally augments it by finding witness points for missing support α vectors and adding the corresponding support α vectors to Ω_{a_t}. Littman, et al. proposed to find a witness point by solving $LP(\alpha, \beta, \Omega_{a_t})$ formulated as follows:

$$
\begin{aligned}
\text{Max} \quad & z(\pi) = \sum_{s \in S} \pi(s)\alpha(s) - \sum_{s \in S} \pi(s)\beta(s)) \\
\text{subject to} \quad & \sum_{s \in S} \pi(s)\alpha(s) \geq \sum_{s \in S} \pi(s)\alpha_i(s), \forall \alpha_i \in \Omega_{a_t} \\
& \sum_{s \in S} \pi(s) = 1 \\
& 1 \geq \pi(s) \geq 0, \forall s \in S,
\end{aligned}
\tag{2.10}
$$

where β is a vector in Ω_{a_t}, and α is constructed based on β. If the optimal objective value $z(\pi^*) > 0$, then the optimal solution, π^*, is a witness point of some missing α-vector. Otherwise, π^* is not a witness point. The witness algorithm is also exponential in the worst case, however it is typically much faster than Sondik's one-pass algorithm in practice [43].

Zhang and Liu (1996) [73] proposed an algorithm, called *incremental pruning*, to solve POMDPs by combining Monahan's method and the witness algorithm. Like the witness algorithm, it represents the conditional value function $v_t(\cdot \mid a_t)$ for each action $a_t \in A_t$ and then focuses on the conditional value function $v_t(\cdot \mid a_t, \theta_t)$ for each observation $\theta_t \in \Theta$ individually, where

$$v_t(\pi_t|\theta_t, a_t) = r_t(\pi_t, a_t)/|\Theta| + \lambda v_{t+1}(\pi_{t+1})\overline{p}_t(\theta_t|\pi_t, a_t),$$

$$v_t(\pi_t|a_t) = \sum_{\theta_t \in \Theta} v_t(\pi_t|\theta_t, a_t),$$

and

$$v_t(\pi_t) = \max_{a_t \in A_t} v_t(\pi_t|a_t).$$

For each action a_t and observation θ_t, it first generates an α-vector set $\Omega_{a_t}^{\theta_t}$ to represent $v_t(\cdot \mid a_t, \theta_t)$ from the previous α-vector set Ω_{t+1} as below:

$$\Omega_{a_t}^{\theta_t} = \{\alpha_t = r_t(a_t)/|\Theta| + \lambda Q_t(\theta_t, a_t)P_t(\theta_t)\alpha_{t+1} \mid \forall \alpha_{t+1} \in \Omega_{t+1}\}.$$

$\Omega_{a_t}^{\theta_t}$ is pruned to a minimal set that includes only support α-vectors of $v_t(\cdot \mid a_t, \theta_t)$ by solving LPs using Monahan's method. Next, an α-vector set Ω_{a_t} for representing $v_t(\cdot \mid a_t)$ can be constructed as follows:

$$\Omega_{a_t} = (\bigoplus_{\theta_{t+1} \in \Theta} \Omega_{a_t}^{\theta_t}(x)),$$

where the $\bigoplus$ is the *cross sum* operation of two sets of vectors defined as:

$$A \bigoplus B = \{\alpha + \beta | \alpha \in A, \beta \in B\}.$$

Next, Ω_{a_t} is pruned using Monahan's method. Finally, Ω_t is constructed as

$$\Omega_t = \bigcup_{a_t \in A} \Omega_{a_t},$$

and it is pruned to a minimal size to represent $v_t(\cdot)$. Incremental pruning has been shown to be faster than other exact algorithms including the witness algorithm [8].

All of the above exact algorithms are intractable for large POMDPs, therefore a number of approximate algorithms have been proposed. Lovejoy [46] categorized the approximations into two categories: *finite-memory* and *finite-grid*. In finite-memory, such as the *policy improvement algorithm* proposed by Howard (1973) [32], only a finite number of decision policies are kept at each epoch, in other words, the α-vector set, Ω_t, is truncated to a finite set at each epoch. Therefore the value function, $v_t(\cdot)$, is approximated by the finite α-vector set. In finite-grid approximations, such as the *fixed-grid algorithm* proposed by Eckles (1966) [17], the continuous belief space π is descritised as a finite number of grids $\pi_n, n = 1, ..., N$. Thus the value function, $v_t(\pi)$, is approximated by $\{v_t(\pi_n) \mid n = 1, ..., N\}$. A more through review of these types of methods can be found in [46].

2.5 Applications of MDPs and POMDPs in Medical Decision Making

MDPs are a powerful and appropriate technique for many medical treatment decisions. MDPs provide optimal policies to stochastic and dynamic decisions. Examples of such decisions naturally arise in finding optimal disease treatment plans. For example, Schaefer, et al. (2005) summarized some of the most successful applications of MDPs to medical treatment decisions [59]. For example, Ahn and Hornberger (1996) used an MDP model for a kidney transplantation problem in which patients may accept or reject an offered kidney based on the quality of the organ [2]. As another example, Alagoz et al. (2004) [3] used an MDP model for deciding the optimal time to perform a living-donor liver transplantation.

POMDPs fit the structures of medical decisions where a patient's true health status is not known with certainty and can only be inferred probabilistically based on imperfect medical tests. In such cases, a POMDP represents the decision making process more accurately than an MDP. Smallwood et al. (1971) [63] was among the first to suggest formulating medical decision problems in the POMDP framework. Since then, there have been many successful POMDP

applications to medicine and healthcare [9]. Following are some examples.

Hauskrecht and Fraser (2000) [25] applied a POMDP formulation to the problem of treating patients with ischemic heart disease (IHD). The state of the patient was described by a variety of variables including the level of coronary heart disease, ischemia level, history of coronary artery bypass grafting, history of percutanerous transluminal coronary angioplasty, and stress test results. The uncertainty of the patient health state arises from the inability to know exactly the level of coronary artery occlusion or the homodynamic impact of that occlusion on myocardial ischemia. Some variables, such as level of chest pain, are directly observable. The authors framed their POMDP as an infinite-horizon discounted model that seeks a treatment policy that minimizes total lifetime costs. Their POMDP was solved with approximating methods very efficiently in generating good treatment strategies for IHD.

Maillart, et al. (2008) [49] used a POMDP to develop a cost benefit analysis of mammogram frequency and treatment options for breast cancer. The part of the model that was partially observable was the patient's cancer state. The goal was to minimize the total expected cost over a patient's lifetime, where costs were based on the patient's condition, exams, and treatment options. Ivy solved this POMDP problem and characterized optimal decision regions based on the perceived probabilities of the different states of breast cancer.

Leshno, et al. (2003) [42] conducted a cost-effective analysis of screening for Colorectal cancer (CRC) using a POMDP model. Screening policies were evaluated using simulation based on the POMDP model. Their study revealed that it is highly cost-effective to screen average-risk asymptomatic individuals.

Zhang, et al. (2009) [72] [71] used a POMDP model to estimate the benefit of PSA-based screening for prostate cancer. Zhang first solved one POMDP to maximize individual's expected quality-adjusted life years (QALYs), then he solved another POMDP to maximize the expected monetary value based on societal willingness to pay for QALYs and the cost of PSA testing, prostate biopsies and treatments. An age and belief dependent biopsy referral threshold is calculated with a xed-nite-grid approximation algorithm. He also proved a number of structural properties of the POMDP including the existence of a control-limit type policy for the biopsy referral decision.

2.6 Contributions of this Thesis to the Literature

This thesis presents several new model formulations for the evaluation of surveillance policies for low risk bladder cancer patients. To our knowledge, we present the first models to study bladder cancer surveillance in low risk patients. We use a partially observable simulation model to compare published international guidelines for bladder cancer surveillance. We extend our simulation model to a POMDP model to determine the optimal surveillance policy that

maximizes a patient's expected QALYs. Based on computational results we develop a number of new insights about the optimal structure of surveillance policies, and the influence of risk factors and gender. We further extend the POMDP model of Chapter 4 to a new POMDP model that involves decisions about the optimal selection and timing of a diagnostic test with imperfect sensitivity and specificity to detect bladder cancer. In Chapter 5 we study the optimal policies based on an additional surveillance protocol that includes a urine based biomarker test. We further investigate easy-to-implement age dependent surveillance schedules and compare them to the optimal surveillance policies. We evaluate the urine based biomarker test in terms of expected QALYs, expected number of cystoscopies, expected number of biomarker tests, by comparing the outcomes from the optimal policy using biomarkers and the one using cystoscopy alone. Finally, we evaluate the incremental benefit of using the FDA approved biomarkers in directing the frequency of diagnostic tests by comparing easy-to-implement heuristic policies using a biomarker with the easy-to-implement heuristic policy using cystoscopy alone.

Chapter 3

Low Risk Bladder Cancer Surveillance Strategies

3.1 Introduction

Carcinoma of the urinary bladder ranks fifth among malignancies, with greater than 70,000 new cases estimated and over 500,000 survivors in 2009 [34]. Typical of epithelial malignancies, bladder cancer incidence is highest in the elderly. Therefore the changing age structure of the U.S. population suggests that the burden of this disease will increase in the future. Clinically, bladder cancer cases are risk-stratified on the basis of stage and grade. The natural history and molecular biology of different risk groups are sufficiently different to suggest the existence of at least three discrete phenotypes: high grade muscle-invasive, and high and low grade non-muscle-invasive [70] [13]. Muscle-invasive disease (stage T2 or greater) accounts for approximately 25% of incident cases, with high risks of metastasis-related morbidity and disease-specific mortality despite radical surgical therapy and systemic chemotherapy. The overwhelming majority of incident bladder cancer presents at a stage superficial to the muscularis propria, broadly defined as non-muscle-invasive bladder cancer (NMIBC). For these NMIBC cases, standard clinical management includes endoscopic transurethral resection of the bladder tumor (TURBT), followed by frequent, invasive surveillance with cystoscopy.

The proximate outcomes for patients with a history of NMIBC are recurrence of NMIBC or progression, defined as recurrence of a tumor with invasion into the muscle. The risks of these distinct outcomes differ starkly between high and low grade NMIBC. Low grade noninvasive tumors account for approximately 70% of all incident NMIBC, and given the comparatively indolent natural history of these cases, likely represent the majority of prevalent bladder cancer cases [7]. In this setting, the principal risk (approximately 40-50% by 5 years [66]) is recurrence of low grade noninvasive tumors, whereas the long-term risk of progression to muscle-invasive

disease is less than 5% and in some series, less than 1% [66] [28] [30]. The long-term progression rates from low grade noninvasive urothelial carcinoma to invasive cancer, in some respects, parallel those of colorectal adenomata [50]. In contrast, high grade NMIBC cases have not only higher rates of recurrence, typically of high grade NMIBC, but more importantly, substantially higher risks of progression to muscle-invasive disease, up to 50-75% by 5 years [66] [16]. Given that low grade noninvasive bladder cancer accounts for nearly half of the overall incident cases, these differences in phenotype-associated outcomes argue for consideration of a risk-adjusted approach to surveillance.

Cystoscopy is the reference standard for surveillance of patients with a history of NMIBC. In the context of the heterogeneous natural history of NMIBC, there is an interesting and substantial difference between the recommendations of the relevant European and U.S. clinical practice guidelines. The European Association of Urology (EAU 2009) [4] advocates explicit risk stratification and, among low risk cases, recommends surveillance cystoscopy at 3 months, 9 months, and annually thereafter for patients without recurrence. In contrast, the American Urological Association (AUA) guidelines (e.g. AUA 2007 [24]) do not explicitly risk-stratify surveillance recommendations, and outline a schedule of cystoscopy every 3 months for 2 years, every 6 months for the next 3 years, and annually thereafter for patients without recurrence. The AUA guidelines acknowledge the potential appropriateness of less intensive regimens for select patients, but no explicit guidance is given, and a one size fits all, relatively intensive approach would be consistent with the AUA guideline recommendations.

For NMIBC patients, surveillance policies must trade off the benefit of early detection of recurrence and/or progression against the economic and quality of life costs of frequent, invasive surveillance. Cystoscopy can be painful and anxiety-provoking for patients [5] [29]. Given the variable natural history of NMIBC, these tradeoffs can differ greatly not only in terms of cancer-specific risks but also, with the predominantly elderly demography of the bladder cancer population, in terms of age and associated competing risks to survival.

We developed a partially observable Markov model to compare surveillance strategies for patients with low grade noninvasive bladder cancer. We evaluated strategies based on QALYs; we also performed a bicriteria analysis to compare expected life-long progression rate versus the number of cystoscopies. We found that the best strategy is sensitive to the disutility of cystoscopy, age, and all other cause mortality. We conclude that the best surveillance strategy is highly dependent on the individual patient. The lower the disutility of cystoscopy, the more frequently cystoscopy should be performed. Older patients and patients with comorbidity should be screened less frequently.

3.2 Model Formulation

The EAU and AUA guidelines represent the reference standard practice guidelines for bladder cancer surveillance in Europe and the U.S., respectively, as summarized in Table 3.1. For low risk patients, the EAU suggests cystoscopy at 3 months, if negative then follow-up cystoscopy is advised at 9 months, and subsequently at yearly intervals for 5 years. The AUA guidelines do not make specific, explicit recommendations for low risk disease; instead they recommend a more intense surveillance schedule for all patients, regardless of risk stratum. The AUA guidelines do suggest, however, consideration of less intensive regimens (not further specified) based on individual patient factors. In this context, we also considered additional hypothetical dynamic strategies.

Table 3.1: Published guidelines for surveillance of low risk bladder cancer patients.

Guidelines	**Recommendations for Low Risk Bladder Cancer Patients**
EAU [4]	Cystoscopy at 3 months; If negative, next cystoscopy at 9 months; If negative, cystoscopy yearly for 5 years.
AUA [24]	No low-risk-stratum-specific schedule is explicitly advocated; the following is mentioned: Every 3 months in the first two years; Every 6 months for subsequent 2-3 years; Annually thereafter.

Since the guidelines advocate cystoscopy every 3 months, with increasing intervals if appropriate, we also evaluated additional strategies with increasing intervals for patients who do not have a recurrence. Since it is a standard among published guidelines to do the first cystoscopy at 3 months, and to stop surveillance for low risk patients after 5 years, all strategies we evaluated discontinued surveillance at 5 years provided the patient had no recurrence during that period. With this in mind, we evaluated a series of dynamic strategies, denoted by D_i, in which cystoscopies are performed at increasing intervals of $3, 3+i, 3+2i, 3+3i, ...$ up to 5 years if no recurrence occurs. Thus, for example, strategy D_3 involved cystoscopies at month $0, 3, 9, 18, 30, 45$.

3.2.1 Partially Observable Markov Model

The patient's health state at any given decision epoch is not known with certainty. Observations are obtained as a result of cystoscopy and a positive observation will trigger treatment. Therefore, frequent cystoscopy will result in a higher probability of diagnosing recurrent tumors

prior to progression. On the other hand, more frequent surveillance results in a reduction in expected rewards due to the disutility associated with cystoscopy. To compare how strategies balance these competing factors we used a Markov model comprised of health states that define the natural history of bladder cancer, treatment, and death from bladder cancer and all other causes. The model is illustrated in Figure 3.1. The model formulation is defined as follows:

Table 3.2: Markov states of the natural history model and treatment model for bladder cancer. Note that the stratification of low risk, intermediate risk and high risk is with regard to nonmuscle invasive bladder cancer, and the high risk state is actually lower risk than muscle-invasive bladder cancer state.

Index	**Natural History States**	**Description**
1	Low Risk Disease Free	History of small volume, low-grade Ta
2	Intermediate Risk NMIBC	Recurrent low-grade Ta or multi-focal and/or large volume low-grade Ta
4	High Risk NMIBC	High-grade Ta, T1, and/or CIS
6	Muscle Invasive Bladder Cancer	T2, T3, T4 tumors
7	Death from other causes	Competing Mortality
Index	**Treatment States**	**Description**
3	Intermediate Risk Disease Free Following Treatment	Intermediate Risk Bladder Cancer Treated, Disease Free
5	High Risk Disease Free Following Treatment	High Risk Bladder Cancer Treated, Disease Free

Decision epochs: We let $t = 1, 2, ..., T$ index monthly decision epochs over the course of a bladder cancer patient's lifetime, where T represents a reasonable upper limit on a patient's age (e.g. 100 years).

States: The patients' health state at epoch t is indexed by $s_t \in S$, where $S = \{1, ..., H, H+1\}$. State H represents the state of muscle invasive bladder cancer and $H+1$ represents the state of death. States 1 to $H-1$ were developed from the EAU classification of non-muscle invasive patients, based on prognostic factors including tumor stage, tumor grade, tumor size, and recurrence rate [4]. A descriptive list of states including the index for each state is provided in Table 3.2. The model includes the following five natural history health states: low risk NMIBC, intermediate risk NMIBC, high risk NMIBC, muscle invasive bladder cancer, and death from bladder cancer and all other causes. For high risk and muscle invasive disease, cystectomy is a common recommendation [18]. For small, recurrent, low grade bladder tumors, some studies suggest it may not be necessary to remove tumors promptly at recurrence [64] [57]

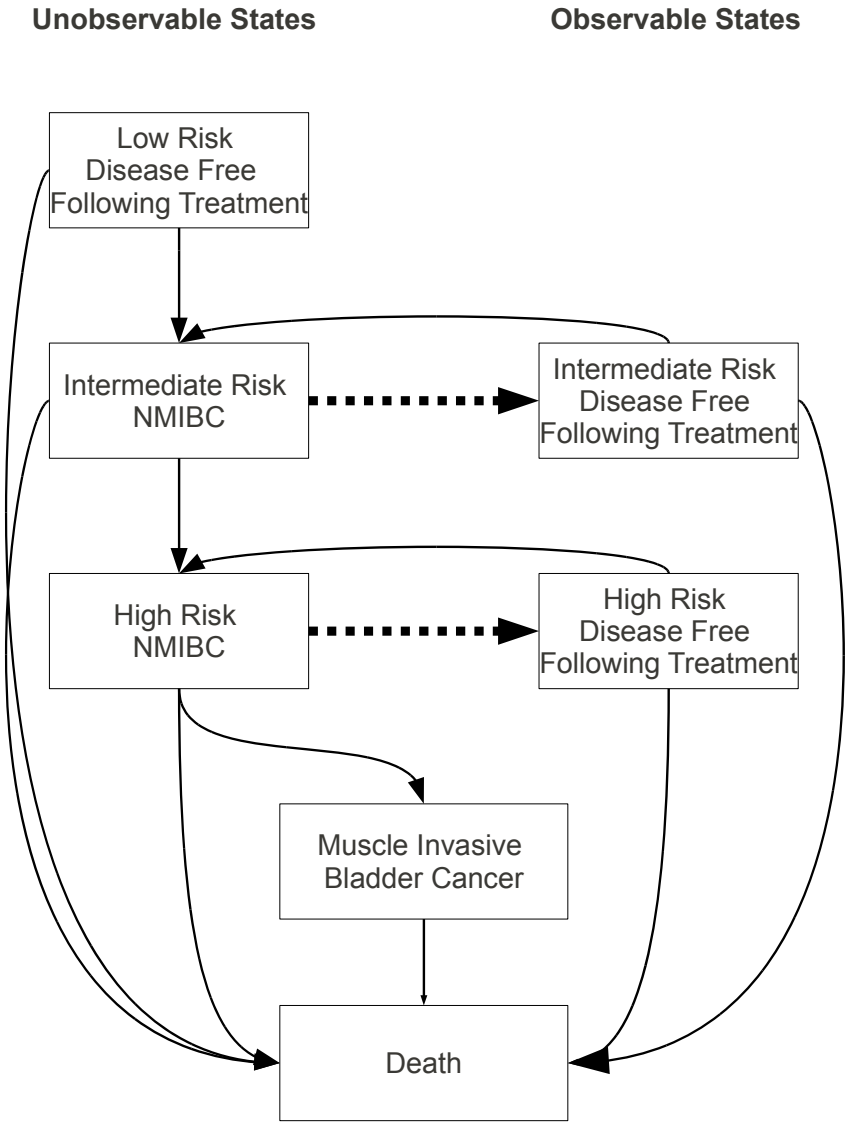

Figure 3.1: The states and possible transitions between states for a patient with bladder cancer. The solid lines indicate the probabilistic transitions. The dashed lines indicate the transitions resulting from the detection of bladder cancer via cystoscopy.

[21]. In our model low risk patients continue to follow cystoscopic surveillance, and treatment is triggered when the patient progresses to intermediate or high risk. We include the following two treatment states: (1) intermediate risk disease free following treatment, and (2) high risk disease free following treatment. Patients are only low risk until the first occurrence, after which they are intermediate risk (recurrent Low grade Ta) or high risk (high grade recurrence).

Transition Probabilities: We let $p_{s_t,s_{t+1}}(t)$ denote the probability that the patient will be in health state s_{t+1} in epoch $t+1$ given he is in state s_t in epoch t. The transition probability matrix, $P(t)$, can be written as follows (blank spaces indicate zeros):

$$P(t) = \begin{bmatrix} p_{1,1}(t) & p_{1,2}(t) & & & & & p_{1,7}(t) \\ & p_{2,2}(t) & & p_{2,4}(t) & & & p_{2,7}(t) \\ & p_{3,2}(t) & p_{3,3}(t) & & & & p_{3,7}(t) \\ & & & p_{4,4}(t) & & p_{4,6}(t) & p_{4,7}(t) \\ & & & p_{5,4}(t) & p_{5,5}(t) & & p_{5,7}(t) \\ & & & & & p_{6,6}(t) & p_{6,7}(t) \\ & & & & & & p_{7,7}(t) \end{bmatrix}. \tag{3.1}$$

We define parameters for annual rates of bladder cancer mortality, bladder cancer recurrence and progression, and all other cause mortality in Table 3.3. At each health state, patients may die from other causes; we assume $p_{1,7}(t) = p_{2,7}(t) = p_{3,7}(t) = p_{4,7}(t) = p_{5,7}(t) = \delta(t)$. In the MIBC state, patients may die from BC with probability δ_{BC}, or other causes with probability $\delta(t)$. Therefore we assume $p_{6,7}(t) = \delta_{BC} + \delta(t)$. In the LRDF state, patients may have recurrent cancer of intermediate risk with probability $p_{1,2}(t) = \tau_{LR}$. In the IRBC state, bladder cancer may progress to the HRBC state with probability τ_{IR}, thus $p_{2,4}(t) = \tau_{IR}$. In the IRDF state, patients may have recurrence with probability γ_{IR}, thus $p_{3,2}(t) = \gamma_{IR}$. Similarly, in the HRDF state, patients may have recurrence with probability γ_{HR}, thus $p_{5,4}(t) = \gamma_{HR}$. In the HRBC state, patients may progress to the MIBC state with probability τ_{HR}, thus $p_{4,6}(t) = \tau_{HR}$. Finally, the Death state is an absorbing state, with $p_{7,7}(t) = 1$.

Decision: The cystoscopy surveillance decision at epoch t is indexed by $a_t \in A_t = \{$Cystoscopy (C), No Cystoscopy (N)$\}$.

Rewards: The rewards for state s_t and action a_t are denoted by $r(s_t, a_t)$. They are measured in QALYs, by subtracting the disutilities of cystoscopy and treatment associated with decision a_t. The disutilities are defined in Table 3.4.

We use $R(a_t) = \{r(s_t, a_t)\}$ to denote the *reward vector* which can be written as follows:

Table 3.3: Model parameters and data sources for monthly mortality, bladder cancer mortality, bladder caner recurrence and progression rates.

Parameter	**Description**	**Data Sources**
$\delta(t)$	Mortality rate at age t	CDC [26]
δ_{BC}	Bladder cancer mortality	Madersbacher [48]
φ_C	Sensitivity of cystoscopy	Grossman [22]
γ_{IR}	Recurrence rate of intermediate risk NMIBC after treatment	EORTC [14]
γ_{HR}	Recurrence rate of high risk BC after treatment	
ρ_{LR}	Progression rate of low risk BC after treatment	
ρ_{IR}	Progression rate of intermediate risk NMIBC after treatment	
ρ_{HR}	Progression rate of high risk NMIBC after treatment	
τ_{LR}	Probability of transition from LRDF to IRBC	
τ_{IR}	Probability of transition from IRBC to HRBC	
τ_{HR}	Probability of transition from HRBC to MIBC	

Table 3.4: Model parameters and data sources for utilities and disutilities for estimating QALYs. The base case values are drawn from Kulkarni's study [41].

Parameter	**Description**	**Value**
μ_C	Disutility of Cystoscopy	0.0025
μ_T	Disutility of TURBT	0.03
μ_{Chemo}	Disutility of Chemotherapy	0.02
μ_{BCG}	Disutility of BCG Maintenance	0.09
r_{LRDF}	Utility in State Low Risk Disease Free Following Treatment	0.98
r_{IRBC}	Utility in State Intermediate Risk NMIBC	0.95
r_{IRDF}	Utility in State Intermediate Risk Disease Free Following Treatment	0.95
r_{HRBC}	Utility in State High Risk NMIBC	0.93
r_{HRDF}	Utility in State High Risk Disease Free Following Treatment	0.93
r_{MIBC}	Utility in State High Muscle Invasive Bladder Cancer	0.80

$$R(N) = \begin{bmatrix} r_{LRDF} \\ r_{IRBC} \\ r_{IRDF} \\ r_{LRBC} \\ r_{HRDF} \\ r_{MIBC} \\ 0 \end{bmatrix}, \qquad R(C) = \begin{bmatrix} r_{LRDF} - \mu_C \\ r_{IRBC} - \mu_C - \mu_T - \mu_{Chemo} - \mu_{BCG} \\ r_{IRDF} - \mu_C \\ r_{HRBC} - \mu_C - \mu_T - \mu_{Chemo} - \mu_{BCG} \\ r_{HRDF} - \mu_C \\ r_{MIBC} \\ 0 \end{bmatrix}.$$

A surveillance strategy defines the sequence of decisions, $\xi = \{a_1, a_2, ..., a_T\}$, about whether to perform a cystoscopy at each decision epoch, t. To compare strategies we estimated the total expected QALYs, $E_\xi[\sum_{t=1}^{T} r(s_t, a_t)]$, over the patient's lifetime. We also estimated expected life-long progression rate and the expected number of cystoscopies.

3.3 Data Sources

Transition probabilities are derived from the EORTC risk table [14], CDC mortality table [26], and survival data [61], summarized in Table 3.3. The EORTC risk tables were developed from pooled individual patient-level data from 2596 patients with NMIBC enrolled in 7 clinical trials. Annual (1 through 5-year) probability estimates of NMIBC recurrence and progression to muscle-invasive bladder cancer are calculated on the basis of coefficients from clinicopathological variables in multivariate logistic regression models. A recent study [33] of 13 cancers (not including bladder cancer) provides evidence that conditional survival rate increase with the time since diagnosis of cancer. To incorporate this into the bladder cancer survival probability, we used a yearly discounting factor, γ, and we assumed the recurrence rate in year t, conditional on remaining disease free, is $p_t = p_1\gamma^{t-1}$. We used 1 year and 5 year recurrence and progression rates from the EORTC table to estimate γ.

It was not possible to estimate the parameters for grade progression, $\tau_{LR}, \tau_{IR}, \tau_{HR}$, directly from the literature. Therefore we estimated them by comparing the model outputs with published progression rates $\rho_{LR}, \rho_{IR}, \rho_{HR}$. We denote the model output of the 5 year progression rate of HRBC patients starting at age t as $f_{HR}(\tau_{HR})$. Similarly, we denote the model output of the 5 year progression rate of IRBC at age t as $f_{IR}(\tau_{IR}, \tau_{HR})$, and for LRBC as $f_{LR}(\tau_{LR}, \tau_{IR}, \tau_{HR})$. We estimated τ_{HR}, τ_{IR}, τ_{LR}, as the choices that minimize $\mid f_{HR}(\tau_{HR}) - \rho_{HR} \mid$, $\mid f_{IR}(\tau_{IR}, \tau_{HR}) - \rho_{IR} \mid$, and $\mid f_{LR}(\tau_{LR}, \tau_{IR}, \tau_{HR}) - \rho_{LR} \mid$, respectively.

We estimated the mortality rate of MIBC from survival data of MIBC patients who underwent radical cystectomy. We performed a PubMed search on the recent published literature on the bladder cancer survival from 2000 to 2010 using the following keywords: (bladder cancer [Title/Abstract]) AND survival[Title]) AND radical cystectomy[Title])). We excluded studies

that were not based on patient cohort data or clinical trials, leaving 8 studies in total [15], [47], [65], [61], [48], [39], [19], [11], as summarized in Table 3.5. We performed sensitivity analysis with respect to bladder cancer mortality estimated from these studies.

Table 3.5: Characteristics of papers studying survival of muscle invasive bladder cancer patients. MI = muscle invasive. DSS = disease-specific survival. RC = radical cystectomy. Chemo = chemotherapy. RR = radical radiology. RFS = recurrence-free survival. OS = overall survival. PMI = primary muscle invasive. PRMI = progressive muscle invasive. C.I.=95% confidence interval.

Authors	Vries [15]	Lund [47]	Stein [65]	Shariat [61]	Madersbacher [48]	Koga [39]	Ferreira [19]	Chahal [11]
Published Year	2010	2010	2001	2006	2003	2008	2007	2003
Country	Netherlands	Denmark	U.S.	U.S.	Switzerland	Japan	Brazil	UK
Study period	1987-2005	1996-2007	1971-1997	1984-2003	1985-2000	1997-2006	1993-2005	1993-1996
Sample size	188	3997	1054	888	507	97	242	398
Median age	61	72	66		66	70		69
Male	75%	74%	80%		79%	73%		60%
Comorbidity		43%	28%					
Treatment	RC		RC		RC	RC, Chemo	RC	RR,RC
5y RFS(C.I.)			68%	58% (56-60)	62%			
5y OS(C.I.)		45% (41-48)	66%		59%	66%		37%
5y DSS(C.I.)				66% (64-68)		74%		
PMI 5y DSS(C.I.)	49% (40-60)						52%	
PRMI 5y DSS(C.I.)	52% (37-74)						58%	

The base case mortality rate, δ_{BC}, was estimated from the 5-year disease specific survival of MIBC patients reported by Shariat (2006) [61] because this study had the largest sample size for bladder cancer disease specific survival. The authors reported results based on a multi-institutional database consisting of 888 consecutive patients with bladder transitional cell carcinoma who were treated with radical cystectomy and pelvic lymphadenectomy at 3 academic centers in the U.S. between 1984 and 2003. Mortality rates from all other causes were estimated from the statistics reports published by the US Centers for Disease Control (CDC) (2009) [26]. We transformed yearly rates, denoted by p_y, to corresponding monthly rates, p_m, by the formula $(1-p_m)^{12} = (1-p_y)$. The base case of cystosocpy sensitivity is set to be 0.95 [22].

We used estimates of utilities and disutilities reported by Kulkarni et al. (2007) [41]. The estimates are summarized in Table 3.4.

3.4 Results

We used our Markov model to compare the EAU, AUA, and dynamic strategies $D_1, D_2, ..., D_{12}$. The base case parameter choices are defined in Tables 3.3 and 3.4. We performed sensitivity analysis with respect to this base case. The results, if not specified otherwise, are based on the base case scenario. We used C++ to implement the simulation process. For each scenario, we used 1,000,000 samples to estimate the mean and 95% confidence intervals. In all the scenarios presented below, the simulation process were be completed within 1 hours on a PC with quad core 2.83GHz CPU and 8GB RAM.

3.4.1 Model Validation

We compared our model estimates of survival to those published in the EUROCARE-3 study (2003) [58], which studied patients diagnosed with bladder cancer during 1990-1994. The EUROCARE-3 study summarized one-year, three-year and five-year survival by age at diagnosis for 21 European countries. The age ranges reported are 45-54, 55-64, 65-74 and 75-99. For our validation results we assigned the value of each age range to the appropriate range in the EUROCARE-3 study. Most of the bladder cancer patients in the EUROCARE-3 study had muscle invasive disease, however the proportion varies considerably from one country to another. Therefore, we compared the EUROCARE-3 study results to the computational lower bounds and upper bounds of our model, defined by the survival of MIBC patients and the survival of LRBC patients, respectively. As expected the 5-year survival results obtained from our model lie between the survival of MIBC patients and that of HRBC patients starting at any age between 50 and 85.

3.4.2 Base Case Results

From 2004-2008, the median age at diagnosis for bladder cancer was 73 years; approximately 9.6% were diagnosed under age 55, and 13.2% above age 85 [1]. Therefore we compare the EAU and AUA guidelines, and strategies D_1 to D_{12}, for the base case of a 73 year old male and female patient. The base case parameter values are provided in Table 3.4. Note that the value for the disutility of cystoscopy, 0.025, is based on the midpoint of the plausible range in Kulkarni et al (2007). From Figure 3.2, the EAU guideline resulted in higher expected QALYs compared to the more intensive AUA guideline (11.05 vs 10.90 QALYs for a 73 year old male), and dynamic strategy D_{12} resulted in the highest QALYs (11.11 QALYs), which dominates

AUA, EAU, D_1, D_2, D_3 and D_4, but not statistically different than the other policies. The results in Figure 3.2 indicate that the ranking of strategies is similar for female patients.

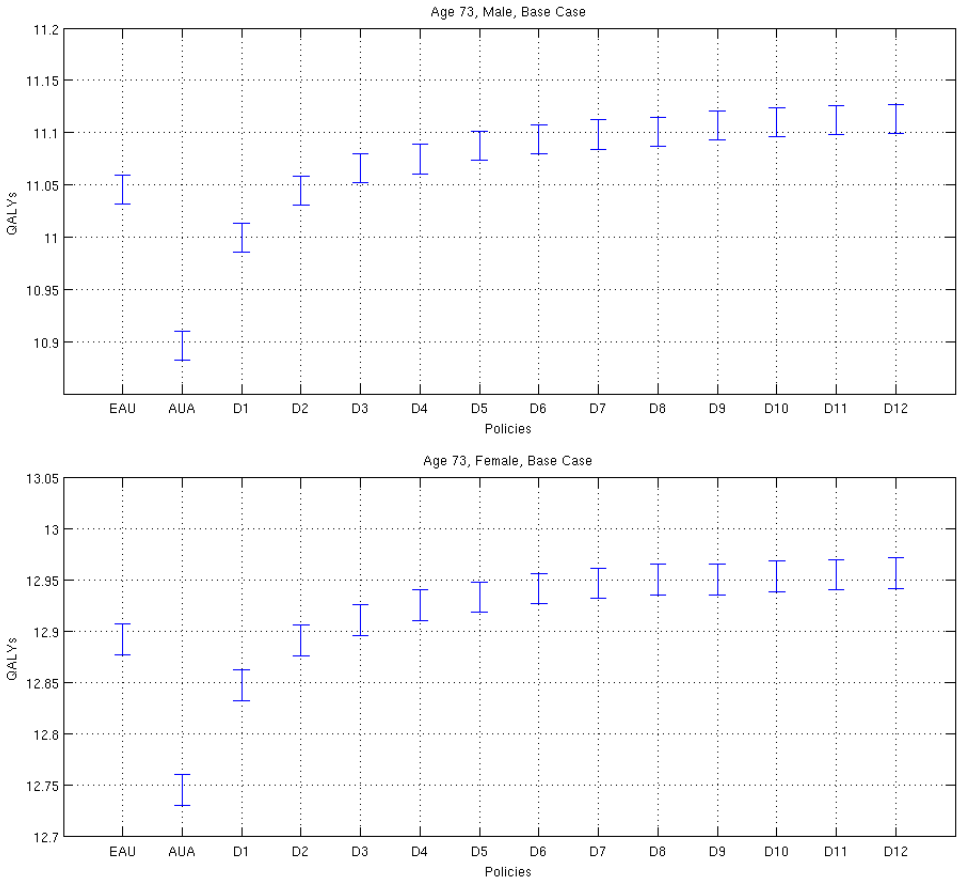

Figure 3.2: Expected QALYs and 95% confidence intervals for all strategies for a 73 year old male and a female patient in the base case.

We compared the surveillance strategies on the basis of expected progression rate versus number of cystoscopies. Figure 3.3 shows the outcome of all strategies for low risk male patients aged 73. The results indicate considerable differences between the EAU and AUA guidelines. For example, the EAU guideline resulted in a higher expected life-long progress rate but with approximately half the number of cystoscopies over the patient's lifetime. The AUA guideline had an absolute reduction of 0.4% in life-long progression rate and increase of 6.48 cystoscopies on average compared to the EAU guideline. No one strategy dominates another; however, there are significant differences in number of cystoscopies relative to changes in expected progression rate.

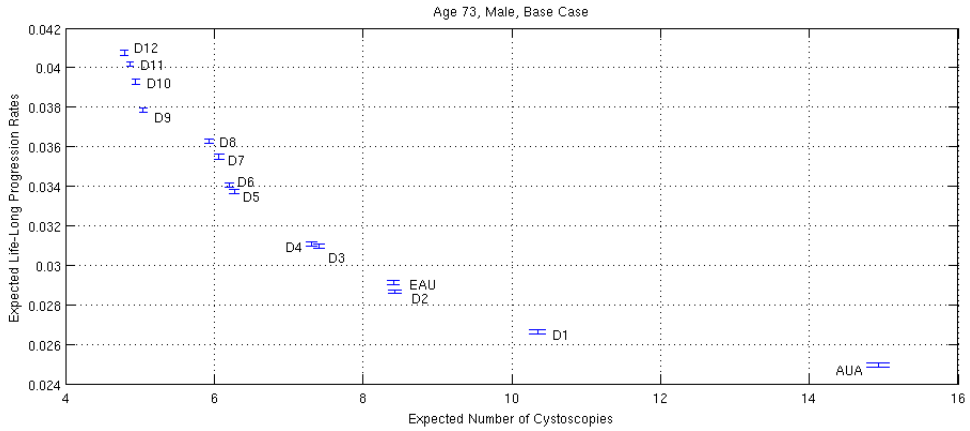

Figure 3.3: The expected life-long progression rate to muscle invasive disease versus number of cystoscopies over a patient's life time for a patient aged 73 under the base case scenario.

3.4.3 Sensitivity Analysis

We performed one-way sensitivity analysis on all strategies for model parameters including disutility of cystoscopy, BC mortality, recurrence and progression rates, and all other cause mortality. The parameter ranges are presented in Table 3.6. In the case of recurrence and progression rates, all parameters were varied simultaneously. Table 3.7 provides the one-way sensitivity analysis for disutility of cystoscopy. The results for disutility of cystoscopy suggest more intensive surveillance when the disutility of cystoscopy is lower. Table 3.8 provides the one-way sensitivity analysis for other cause mortality. Results for all other cause mortality indicate that patients should have less intensive surveillance as all other cause mortality increases. Finally, the ranking of strategies does not change significantly within the range of bladder cancer mortality and recurrence and progression rates, as shown in Tables 3.9 and 3.10.

We evaluated sensitivity of strategies to the starting age of surveillance. We compared the strategies for male patients from 55 through 85 year old, as shown in Table 3.11. The best strategy for age 55 patients is D_8. The best strategy for age 85 patients is D_{12}.

3.5 Discussion

Our model for bladder cancer is based on recent estimates of risk of recurrence and progression derived from the EORTC risk table [14]. The most related work to ours is that of Kent et al [37], [36]. The authors developed a probabilistic model with five health states: free of tumor, with tumor and intact bladder, post-cystectomy and tumor-free, post-cystectomy with tumor, and death. They compared hypothetical surveillance strategies with an optimal strategy that

Table 3.6: The base case values and ranges for the parameters that changed in the one-way sensitivity analysis. (Note that parameter $\delta(t)$ is dependent on age, progression and recurrence rates were varied by a factor of 0.5 and 2.0.)

Description	Parameter	Base Case Value	Lower Bound	Upper Bound
Disutility of Cystoscopy	μ_C	0.025	0.003	0.05
Bladder Cancer Mortality	δ_{BC}	0.011	0.005	0.016
Other Cause Mortality	$\delta(t)$	$\delta(t)$	$0.5 \times \delta(t)$	$2 \times \delta(t)$
Cystoscopy Sensitivity	φ_C	0.95	0.90	1.00
Progression and Recurrence Rates	γ_{IR}	0.030	0.015	0.060
	γ_{HR}	0.075	0.038	0.150
	τ_{LR}	0.002	0.001	0.004
	τ_{IR}	0.008	0.004	0.016
	τ_{HR}	0.070	0.035	0.140

Table 3.7: One-way sensitivity analysis with respect to disutility of cystoscopy on practical, and dynamic policies for 73 year old low risk male patients. The bold font value at each column indicates the best strategy in the corresponding scenario.

Strategies	Disutility of Cystoscopy Lower Bound	Disutility of Cystoscopy Upper Bound
EAU	11.230 (11.216, 11.244)	10.835 (10.821, 10.849)
AUA	11.225 (11.211, 11.239)	10.523 (10.510, 10.537)
D1	11.227 (11.213, 11.241)	10.741 (10.727, 10.755)
D2	11.230 (11.216, 11.244)	10.834 (10.820, 10.848)
D3	11.229 (11.215, 11.243)	10.880 (10.866, 10.894)
D4	11.236 (11.222, 11.250)	10.892 (10.878, 10.906)
D5	11.225 (11.211, 11.239)	10.930 (10.917, 10.944)
D6	11.230 (11.216, 11.244)	10.938 (10.925, 10.952)
D7	**11.232 (11.218, 11.246)**	10.947 (10.933, 10.960)
D8	**11.232 (11.218, 11.246)**	10.953 (10.939, 10.967)
D9	11.218 (11.204, 11.231)	10.980 (10.967, 10.994)
D10	11.219 (11.205, 11.233)	10.987 (10.973, 11.000)
D11	11.219 (11.205, 11.233)	10.990 (10.976, 11.004)
D12	11.219 (11.205, 11.232)	**10.993 (10.979, 11.007)**

Table 3.8: One-way sensitivity analysis with respect to other cause mortality on practical, and dynamic policies for 73 year old low risk male patients. The bold font value at each column indicates the best strategy in the corresponding scenario.

Strategies	Other Cause Mortality Lower Bound	Other Cause Mortality Upper Bound
EAU	16.757 (16.736, 16.779)	8.626 (8.615, 8.637)
AUA	16.627 (16.605, 16.649)	8.476 (8.465, 8.487)
D1	16.711 (16.689, 16.732)	8.580 (8.569, 8.591)
D2	16.752 (16.730, 16.773)	8.624 (8.613, 8.635)
D3	16.773 (16.752, 16.795)	8.646 (8.635, 8.657)
D4	16.808 (16.787, 16.830)	8.653 (8.642, 8.664)
D5	16.789 (16.767, 16.810)	8.670 (8.659, 8.681)
D6	16.814 (16.793, 16.836)	8.674 (8.663, 8.685)
D7	16.826 (16.804, 16.847)	8.677 (8.666, 8.688)
D8	**16.841 (16.819, 16.863)**	8.680 (8.669, 8.691)
D9	16.798 (16.776, 16.819)	8.694 (8.683, 8.705)
D10	16.802 (16.780, 16.824)	8.697 (8.686, 8.708)
D11	16.813 (16.791, 16.834)	8.698 (8.687, 8.709)
D12	16.818 (16.796, 16.840)	**8.700 (8.689, 8.711)**

Table 3.9: One-way sensitivity analysis with respect to bladder cancer mortality on practical, and dynamic policies for 73 year old low risk male patients. The bold font value at each column indicates the best strategy in the corresponding scenario.

Strategies	Bladder Cancer Mortality Lower Bound	Bladder Cancer Mortality Upper Bound
EAU	11.110 (1.096, 11.124)	11.009 (10.995, 11.023)
AUA	10.956 (10.942, 10.970)	10.863 (10.850, 10.877)
D1	11.063 (11.049, 11.077)	10.964 (10.950, 10.978)
D2	11.109 (11.095, 11.123)	11.008 (10.994, 11.021)
D3	11.132 (11.118, 11.146)	11.028 (11.014, 11.042)
D4	11.137 (11.123, 11.151)	11.040 (11.026, 11.053)
D5	11.156 (11.142, 11.170)	11.049 (11.035, 11.063)
D6	11.160 (11.146, 11.174)	11.056 (11.042, 11.070)
D7	11.163 (11.149, 11.177)	11.061 (11.047, 11.075)
D8	11.166 (11.152, 11.180)	11.064 (11.051, 11.078)
D9	11.180 (11.166, 11.194)	11.066 (11.052, 11.079)
D10	11.183 (11.169, 11.197)	11.070 (11.056, 11.084)
D11	11.184 (11.170, 11.198)	11.072 (11.058, 11.086)
D12	**11.186 (11.172, 11.200)**	**11.072 (11.059, 11.086)**

Table 3.10: One-way sensitivity analysis with respect to progression and recurrence rates on practical, and dynamic policies for 73 year old low risk male patients. The bold font value at each column indicates the best strategy in the corresponding scenario.

Strategies	Progression and Recurrence Rates Lower Bound	Progression and Recurrence Rates Upper Bound
EAU	11.289 (11.275, 11.304)	10.406 (10.393, 10.419)
AUA	11.132 (11.117, 11.146)	10.292 (10.280, 10.305)
D1	11.242 (11.228, 11.257)	10.373 (10.360, 10.386)
D2	11.289 (11.274, 11.303)	10.405 (10.393, 10.418)
D3	11.312 (11.297, 11.326)	10.419 (10.406, 10.431)
D4	11.317 (11.303, 11.332)	10.425 (10.412, 10.437)
D5	11.336 (11.322, 11.351)	10.429 (10.416, 10.442)
D6	11.340 (11.326, 11.355)	10.431 (10.419, 10.444)
D7	11.342 (11.328, 11.357)	10.433 (10.420, 10.445)
D8	11.346 (11.332, 11.361)	10.432 (10.419, 10.445)
D9	11.360 (11.346, 11.375)	10.433 (10.421, 10.446)
D10	11.363 (11.348, 11.377)	**10.434 (10.422, 10.447)**
D11	11.364 (11.350, 11.378)	10.433 (10.420, 10.445)
D12	**11.365 (11.351, 11.380)**	10.432 (10.419, 10.445)

Table 3.11: Comparison of strategies with respect to changes in starting age of surveillance.

Strategies	55	73	85
AUA	21.713 (21.694, 21.732)	11.045 (11.031, 11.059)	5.913 (5.901, 5.924)
EAU	21.605 (21.585, 21.624)	10.896 (10.883, 10.910)	5.792 (5.780, 5.803)
D_1	21.668 (21.649, 21.687)	10.999 (10.986, 11.013)	5.873 (5.862, 5.885)
D_2	21.705 (21.686, 21.724)	11.044 (11.031, 11.058)	5.910 (5.898, 5.921)
D_3	21.722 (21.703, 21.741)	11.066 (11.052, 11.079)	5.929 (5.918, 5.940)
D_4	21.785 (21.766, 21.804)	11.075 (11.061, 11.089)	5.937 (5.926, 5.949)
D_5	21.733 (21.713, 21.752)	11.087 (11.073, 11.101)	5.949 (5.938, 5.961)
D_6	21.774 (21.755, 21.794)	11.094 (11.080, 11.107)	5.955 (5.943, 5.966)
D_7	21.795 (21.776, 21.814)	11.098 (11.084, 11.112)	5.959 (5.948, 5.971)
D_8	**21.813 (21.794, 21.832)**	11.101 (11.087, 11.115)	5.963 (5.952, 5.975)
D_9	21.729 (21.710, 21.749)	11.107 (11.093, 11.120)	5.969 (5.958, 5.981)
D_{10}	21.738 (21.719, 21.757)	11.110 (11.096, 11.124)	5.972 (5.960, 5.983)
D_{11}	21.752 (21.733, 21.772)	11.112 (11.098, 11.126)	5.974 (5.963, 5.985)
D_{12}	21.762 (21.742, 21.781)	**11.113 (11.099, 11.127)**	**5.977 (5.965, 5.988)**

was calculated using a non-linear optimization model to minimize expected delay of tumor detection. In contrast, our study uses recent data, including the EORTC risk tables, to define the model. We compared current guidelines specifically for low risk bladder cancer patients, which are the focus of our study and represent the majority of patients with bladder cancer. We compared strategies on the basis of QALYs, and we used bicriteria analysis to compare expected life-long progression rate versus expected number of cystoscopies over a patient's lifetime.

Our base case results indicated the EAU guideline is associated with greater QALYs than the AUA guideline. We observed no evidence of benefits from very intensive surveillance strategies. For example, strategies D_3 through D_{12} are similar in QALYs but strategy D_{12} results in significantly fewer cystoscopies on average. The best strategy for a 73 year old male patient was found to be D_9, which has nearly half the expected number of cystoscopies compared to EAU over a patients lifetime. The best strategy for a 55 old patient is D_3 and for a 85 year old patient is D_{12}, suggesting that older patients or patients with higher all other cause mortality should generally undergo less intensive surveillance.

Although the differences among strategies on the basis of QALYs is relatively small, our bicriteria analysis revealed there are significant differences among strategies, specifically in the number of cystoscopies. We observed that the number of cystoscopies over a patient's lifetime ranged from 4.13 for strategy D_{12} to 13.76 for the AUA strategy. We found that no one strategy dominated another, i.e., all strategies were on the *efficient frontier*. We observed that the EAU policy resulted in nearly half of the number of cystoscopies with a relative risk reduction of 17% and an absolute risk reduction of 0.4%. The large variation in the number of cystoscopies with the different surveillance strategies, particularly in the context of the very low background rate of progression to invasive cancer in this population, underscores the importance of understanding the quality of life impact of this management practice on patients.

Based on sensitivity analysis, we found that disutility of cystoscopy affected the selection of the best strategy. Generally speaking, patients should undergo more intensive surveillance as the disutility of cystoscopy is reduced. In the extreme case, where disutility is negligible, patients should be screened at each decision epoch since it maximizes the probability of early detection when the patient has entered the intermediate or high risk NMIBC state. All other cause mortality also affects the selection of the best strategy. Patients with higher all other cause mortality should undergo less intensive surveillance. Factors which were not associated with significant differences included bladder cancer mortality, sensitivity of cystoscopy, and progression rate.

3.6 Conclusions

Current guidelines, such as EAU and AUA, do not distinguish patients on the basis of age, comorbidity, or patient's disutility associated with cystoscopy. Our study suggests that the EAU guideline yields higher expected QALYs but also higher life long progression probabilities than the AUA policy. We found that younger patients should be screened more intensively than older patients. Sensitivity analysis showed that patients should undergo more intensive surveillance if the disutility of cystoscopy is reduced. The analysis also showed that patients should also undergo less intensive surveillance as all other cause mortality increases. Based on these results we conclude that patient specific factors such as the presence of comorbidity, or perception of utility loss from cystoscopy, should play a role in determining the best surveillance strategy.

Our study has some limitations. First, in the absence of estimates from the literature, we calculated the monthly stage progression rates in our natural history model by minimizing the absolute deviations of observed 5 year progression rates from the EORTC risk table with our simulation model. Second, we made several assumptions regarding the timing, adherence, and duration of treatment. We assumed that treatment is immediate and adherence to treatment is perfect; we also assumed that treatments will always be done within one decision epoch (one month), while in reality BCG and chemotherapy may last longer. Finally, the disutilities of cystoscopy and various treatment were drawn from Kulkani's study [40], which were in turn drawn from studies of disutility of other invasive procedures, not specifically cystoscopy. Low risk bladder cancer patients have a very low risk of disease progression and cancer-specific mortality. The results of this study suggest that further research to more completely characterize the quality of life impact of frequent, repetitive, invasive procedures in this context are needed to better inform the understanding of the most efficient and effective surveillance program. Finally, we did not include considerations of costs in this model. Given that bladder cancer has been characterized as having the highest cost per-patient from diagnosis to death of all cancer sites, future research could consider the economic, as well as quality of life, impacts of these common practices in a predominantly elderly population [6].

Chapter 4

A POMDP for Low Risk Bladder Cancer Surveillance

4.1 Introduction

Currently there is disagreement in the medical community about how frequently low risk NMIBC patients should receive cystoscopies. Current guidelines are not consistent in their recommendations [4] [54] [52] [24]. In Chapter 3, Monte-Carlo simulation was used to compare guideline recommendations of the European Association of Urology (EAU, 2009) [4], the American Urological Association (AUA, 2007) [24], and other surveillance strategies. In this chapter we consider the design of an optimal surveillance policy. We extend the simulation model in Chapter 3 to a POMDP.

In our POMDP model the health status of a patient is defined by a set of core states. The core states are not directly observable, but they can be observed with a diagnostic test (cystoscopy). Decisions are defined on the belief about the patient's heath status, i.e, the probability that the patient's cancer has recurrence or progression. At each epoch the decision maker chooses either to perform a cystoscopy or defer the decision to test until the next epoch. The belief is updated at each decision epoch. As in Chapter 3, we assume the decision maker's objective is to maximize the patient's expected quality adjusted life years (QALYs).

We begin by formulating the POMDP and discussing its special structure relative to a general POMDP. Next, we use the *incremental pruning algorithm* [8] to find the optimal epoch at which to perform cystoscopy over the course of a patient's lifetime. Using a base case, similar to that of Chapter 3, we compare the optimal surveillance policy based on our model to those of the published EAU and AUA guidelines. We also conduct sensitivity analysis with respect to model input parameters to determine which parameters most influence the optimal policy. We use the results of our numerical experiments to draw some insights about optimal surveillance

strategies and the potential for improving the existing guidelines.

The remainder of this chapter is organized as follows. In Section 4.2, we describe the POMDP model formulation. Next in Section 4.3 we illustrate the methodologies and algorithms used to solve our model. In Section 4.4 we present the numerical results, and in Section 4.5 we discuss the implication of our results. Finally, the conclusions of our study are summarized in Section 4.6.

4.2 POMDP Model Formulation

The objective in our model is to maximize the expected QALYs for a patient. QALYs are estimated by decrementing a normal life year as a result of various events including (a) performing cystoscopy, (b) treatment upon detection of bladder cancer, and (c) long-term complications resulting from treatment. The optimal policy for cystoscopy must trade off the disutility of cystoscopy with possible benefits of early detection of bladder cancer.

The decision process is illustrated by Figure 4.1. At each decision epoch, a decision is made to either perform a cystoscopy, or to defer the decision to the next decision epoch. If the decision is deferred, then the patient's core state changes according to the core state Markov chain. If the patient undergoes cystoscopy, and a negative result is observed, then the patient's core state also changes according to the core state Markov chain. If a positive cystoscopy result is observed (this means it is a true positive as we assume the specificity of cystoscopy is perfect, as in Chapter 3), we assume that the patient will be treated immediately. As in Chapter 3, TURBT, and BCG [55] are assumed to be the treatment for NMIBC disease based on the study of Kulkarni [41]. After treatment, patients are assumed to follow a standard surveillance guideline for intermediate and high risk patients (we assume this because there is general agreement about this guideline in the medical community [4] [54] [52] [24]). Therefore, the patient will leave the decision process. A post-treatment state can be viewed as an absorbing state in our model, with the reward representing expected future QALYs for entering the post-treatment state given the patient undergoes standard surveillance. We use the simulation model from Chapter 3 to estimate the expected rewards of the post-treatment state. Following is a mathematical description of our POMDP model:

Time Horizon: Cystoscopy is performed monthly from a defined starting age (representing an initial occurrence of low risk bladder cancer) to a reasonable upper bound on life span (e.g. age 100). Decision epochs are indexed by $t = 0, 1, 2, ..., T$.

Actions: Action, $a_t \in A_t = \{C, W\}$ denotes the decision to perform a cystoscopy (C) at epoch t or defer the decision to perform cystoscopy (W) until the next decision epoch $t + 1$.

States: At each decision epoch a patient is in one of several health states including low risk

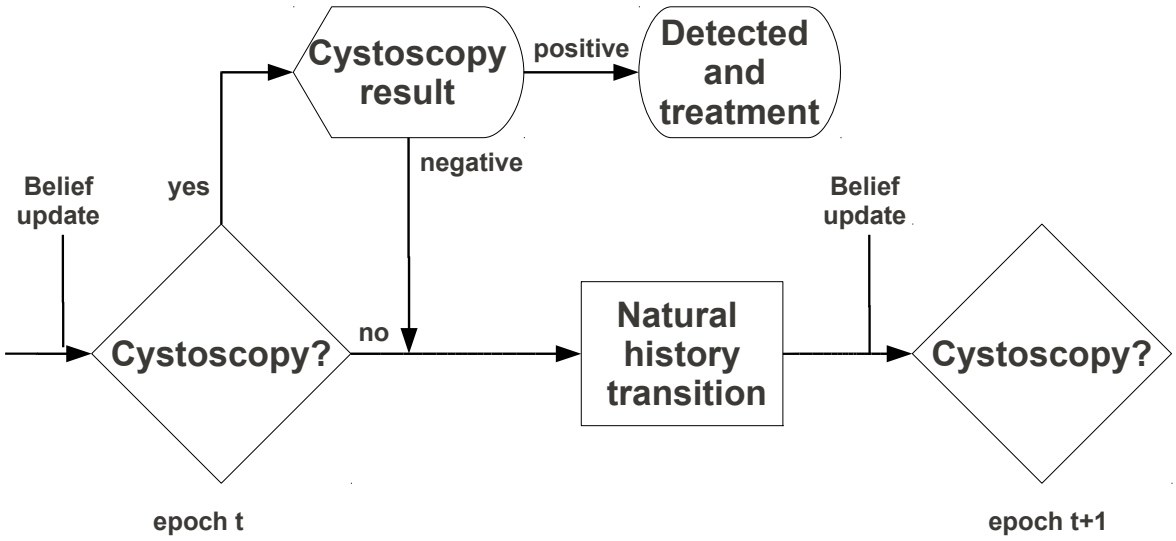

Figure 4.1: Recurring surveillance decision process for low risk bladder cancer patients, in which decision epoches $(t, t+1, ...)$ occur monthly. The belief state is updated before deciding to perform a cystoscopy in each epoch.

disease free following treatment (LRDF), intermediate risk bladder cancer (IRBC), intermediate risk disease free following treatment (IRDF), high risk bladder cancer (HRBC), high risk disease free following treatment (HRDF), muscle invasive bladder cancer (MIBC), and death from bladder cancer and other causes (D).

A descriptive list of all states is provided in Table 3.2. For HRBC and MIBC disease, delay in surgical treatment could result in progression of cancer and decreased survival [18]. Therefore, we assume patients are treated immediately and we represent these treated states, HRDF and MIBC, by absorbing states in our model. As pointed out in Chapter 3, small, recurrent, low grade bladder tumors are slow growing and pose minimal risk; studies indicate that it may not be necessary to remove those tumors promptly at recurrence [64] [57] [21]. Therefore, in our model we assume these low risk patients continue surveillance, and treatment is triggered once a tumor progresses to intermediate risk. Treatment occurs immediately upon detection of intermediate risk recurrent tumors, and the treated state IRDF is also represented by an absorbing state in our model. Figure 4.2 illustrates how treating MIBC, IRDF and HRDF as absorbing states simplifies the model. This simplification does not cause a loss of accuracy in our model since the rewards for states IRDF, HRDF and MIBC are the expected future rewards associated with the underlying Markov reward chain, which is estimated via simulation. We index the possible belief states by s_t where $s_t \in S = \{LRDF, IRBC, HRBC\}$. The set of absorbing state are denoted by $\bar{S} = \{IRDF, HRDF, MIBC, D\}$.

Observations: After a cystoscopy ($a_t = C$) the test result is either positive (p) or negative (n). We let n also denote the observation following the action $a_t = W$. We index the observation by θ_t where $\theta_t \in \Theta = \{p, n\}$.

Information Matrix: Conditional probabilities relate the underlying core states to the observations following a specific action at each decision epoch. We let $q_t(\theta_t|s_t, a_t)$ denote the probability of observing θ_t conditional on being in core state s_t following action a_t at epoch t. We let $Q_t(a_t)$ denote the information matrix conditioning on action a_t, with elements $q_t(\theta_t|s_t, a_t)$. We use f to represent the sensitivity of cystoscopy, resulting in the following two information matrices:

$$Q(C) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \begin{array}{cc} p & n \end{array} \\ \begin{pmatrix} 0 & 1 \\ f & 1-f \\ f & 1-f \end{pmatrix} \end{array}, \qquad Q(W) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \begin{array}{cc} p & n \end{array} \\ \begin{pmatrix} 0 & 1 \\ 0 & 1 \\ 0 & 1 \end{pmatrix} \end{array}.$$

For convenience, we define the follow matrices (note that blank spaces indicate zeros):

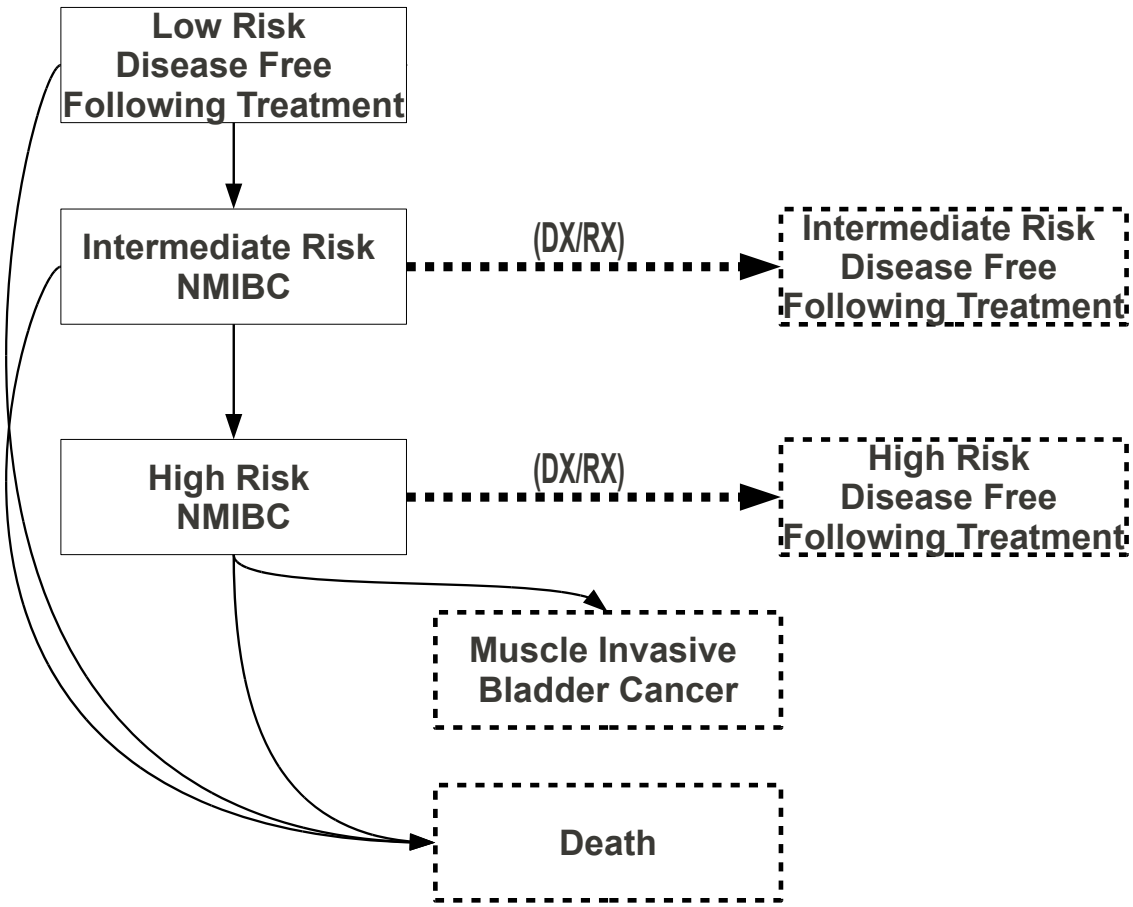

Figure 4.2: The states and possible transitions between states for a patient with bladder cancer in the POMDP model. The solid lines indicate the probabilistic transitions. The dashed lines indicated the transitions resulting from the detection of bladder cancer via cystoscopy. The post-treatment states IRDF, HRDF and MIBC, indicated by dashed boxes, are treated as absorbing states. The rewards for the intermediate and high risk disease free states and the muscle invasive bladder cancer states are estimated using simulation of the underlying Markov reward chain.

$$Q(p,C) = \begin{pmatrix} 0 & & \\ & f & \\ & & f \end{pmatrix}, \qquad Q(n,C) = \begin{pmatrix} 1 & & \\ & 1-f & \\ & & 1-f \end{pmatrix},$$

$$Q(p,W) = \begin{pmatrix} 0 & & \\ & 0 & \\ & & 0 \end{pmatrix}, \qquad Q(n,W) = \begin{pmatrix} 1 & & \\ & 1 & \\ & & 1 \end{pmatrix}.$$

Belief States: We define $\pi_t = (\pi_t(LRDF), \pi_t(IRBC), \pi_t(HRBC))$ as the belief state of the non-absorbing states at epoch t, and $\bar{\pi}_{t+1} = (\bar{\pi}_{t+1}(IRDF), \bar{\pi}_{t+1}(HRDF), \bar{\pi}_{t+1}(MIBC), \bar{\pi}_{t+1}(D))$ as the belief state defined by the absorbing states at epoch $t+1$.

Transition Probabilities: After taking an action a_t, an observation θ_t is observed. If $\theta_t = p$, treatment is triggered and the patient transfers to one of the post-treatment states. If $\theta_t = n$, treatment is not triggered, and the core state transition proceeds. Therefore, the state transition is only dependent on the observation, θ_t, at epoch t. We let $p_t(s_{t+1}|s_t, \theta_t)$ denote the core state transition probability from state s_t to s_{t+1}, given observation θ_t, at epoch t. We let $P_t(\theta_t)$ denote the transition probability matrix conditioned on observation θ_t, with elements $p_t(s_{t+1}|s_t, \theta_t)$ for $s_{t+1} \in S$. Similarly we let $\bar{P}_t(\theta_t)$ denote the transition probability matrix conditioned on observation θ_t, with elements $\bar{p}_t(s_{t+1}|s_t, \theta_t)$ for $s_{t+1} \in \bar{S}$. Therefore, we have the following transition matrices defined as:

$$P_t(n) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \begin{array}{ccc} LRDF & IRBC & HRBC \end{array} \\ \begin{pmatrix} 1-\rho_{LR}-\delta(t) & \rho_{LR} & \\ & 1-\rho_{IR}-\delta(t) & \rho_{IR} \\ & & 1-\rho_{HR}-\delta(t) \end{pmatrix} \end{array},$$

$$\bar{P}_t(n) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \begin{array}{cccc} IRDF & HRDF & MIBC & D \end{array} \\ \begin{pmatrix} & & & \delta(t) \\ & & & \delta(t) \\ & & \rho_{HR} & \delta(t) \end{pmatrix} \end{array},$$

and

$$P_t(p) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \begin{array}{ccc} LRDF & IRBC & HRBC \end{array} \\ \begin{pmatrix} 1-\rho_{LR}-\delta(t) & \rho_{LR} & \\ & & \\ & & \end{pmatrix} \end{array},$$

$$\bar{P}_t(p) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \begin{array}{cccc} IRDF & HRDF & MIBC & D \end{array} \\ \begin{pmatrix} & & & \delta(t) \\ 1-\delta(t) & & & \delta(t) \\ & 1-\delta(t) & & \delta(t) \end{pmatrix} \end{array},$$

where ρ_{LR}, ρ_{IR}, ρ_{HR} are the recurrence rates of low risk, intermediate risk and high risk NMIBC patients, and $\delta(t)$ is other cause mortality, defined in Table 3.3.

Rewards: Performing a cystoscopy imposes a disutility on the patient since it is an invasive test. We let $r_t(s_t, a_t)$ be the core state expected immediate reward (measured in QALYs) given the patient is in state s_t after taking action a_t at time epoch t. After a cystoscopy is performed, the disutility associated with cystoscopy, μ_C, is subtracted; if the patient is in core state IRBC or HRBC, then with probability, f (the sensitivity of cystoscopy), a positive results is observed, after which treatment is triggered and disutilities of treatment, μ_T and μ_{BCG}, is collected. We let $r_t(s_t, \theta_t)$ be the core state immediate reward associated with treatment triggered by observation θ_t. Thus, the belief state expected immediate reward with action a_t can be written as:

$$r_t(W) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \\ \begin{pmatrix} r_{LRDF} \\ r_{IRBC} \\ r_{HRBC} \end{pmatrix} \end{array}, \qquad r_t(I) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \\ \begin{pmatrix} r_{LRDF} - \mu_C \\ r_{IRBC} - \mu_C - f(\mu_T - \mu_{BCG}) \\ r_{HRBC} - \mu_C - f(\mu_T - \mu_{BCG}) \end{pmatrix} \end{array},$$

where the utilities r_{LRDF}, r_{IRBC}, r_{IRDF}, r_{HRBC}, r_{HRDF}, r_{MIBC} and the disutilities of treatment μ_C, μ_T, and μ_{BCG} are as defined in Table 3.4. The expected reward vector of elements for a patient entering the absorbing states, IRDF, HRDF, MIBC, is calculated using the simulation

model of Chapter 3. They are denoted as:

$$R_t = \begin{matrix} IRDF \\ HRDF \\ MIBC \\ D \end{matrix} \begin{pmatrix} R_t(LRDF) \\ R_t(HRDF) \\ R_t(MIBC) \\ 0 \end{pmatrix}.$$

Similarly, we denote the expected reward at LRDF, IRBC, and HRBC at the last decision epoch T as $R_T(LRDF)$, $R_T(IRBC)$, and $R_T(HRBC)$, respectively.

4.2.1 POMDP Model Structure

It is well known that POMDPs can be reduced to an equivalent completely observable Markov decision process on the continuous belief states. The optimal value function can be written as:

$$v_t(\pi_t) = \max_{a_t \in A_t} \left\{ r_t(\pi_t, a_t) + \lambda \sum_{\theta_t \in \Theta} (v_{t+1}(\pi_{t+1}) + \bar{\pi}_{t+1} R_{t+1}) Pr(\theta_t | \pi_t, a_t) \right\}, \forall \pi_t, t = 1, ..., T-1, \tag{4.1}$$

and the value function of the terminal period, T, is defined as:

$$v_T(\pi_T) = \sum_{s_T \in S} \pi_T(s_T) R_T(s_T), \forall \pi_T. \tag{4.2}$$

The optimal action is defined as:

$$a_t^*(\pi_t) = \arg\max_{a_t \in A_t} \left\{ r_t(\pi_t, a_t) + \lambda \sum_{\theta_t \in \Theta} (v_{t+1}(\pi_{t+1}) + \bar{\pi}_{t+1} R_{t+1}) Pr(\theta_t | \pi_t, a_t) \right\}, \forall \pi_t, t = 1, ..., T-1, \tag{4.3}$$

where

$$Pr(\theta_t | \pi_t, a_t) = \sum_{s_t \in S} \pi_t(s_t) q_t(\theta_t | s_t, a_t).$$

Bayesian updates are defined by the following formula:

$$\pi_{t+1}(s_{t+1}) = \frac{\sum_{s_t \in S} \pi_t(s_t) q_t(\theta_t | s_t, a_t) p_t(s_{t+1} | s_t, \theta_t)}{\sum_{s_t \in S} \pi_t(s_t) q_t(\theta_t | s_t, a_t)}, \forall s_{t+1} \in S, \tag{4.4}$$

and

$$\bar{\pi}_{t+1}(s_{t+1}) = \frac{\sum_{s_t \in S} \pi_t(s_t) q_t(\theta_t | s_t, a_t) \bar{p}_t(s_{t+1} | s_t, \theta_t)}{\sum_{s_t \in S} \pi_t(s_t) q_t(\theta_t | s_t, a_t)}, \forall s_{t+1} \in \bar{S}. \tag{4.5}$$

Equations (4.4) and (4.5) provide a means to update the belief state of a patient based on their prior belief state and their most recent action and most recent observation.

4.3 Methodology

Our POMDP is a finite horizon problem with three months as one decision epoch. This is a reasonable assumption, as the published guidelines such as the AUA and EAU guidelines only consider multiples of three months as the surveillance intervals. For such problems it is generally hard to find optimal policies with exact algorithms due to geometric growth in the policy space. However our POMDP has a special structure which significantly reduces its computational complexity such that it can be solved in a reasonable time using incremental pruning. In this section, we first briefly review the incremental pruning algorithm. Next, we describe the special characteristics of our POMDP and illustrate how these can be exploited to reduce its computational complexity.

Incremental pruning uses the conditional value function $v_t(\cdot \mid a_t)$ for each action $a_t \in A_t$, and then focuses on the conditional value function $v_t(\cdot \mid a_t, \theta_t)$ for each observation $\theta_t \in \Theta$ individually. For each action a_t and observation θ_t, it first generates an α-vector set of $\bar{\Omega}_{a_t}^{\theta_t}$ to represent $v_t(\cdot \mid a_t, \theta_t)$ from the previous α-vector set Ω_{t+1}. Then $\bar{\Omega}_{a_t}^{\theta_t}$ is pruned to a minimal set $\Omega_{a_t}^{\theta_t}$ with cardinality less than or equal to that of Ω_{t+1}. Next, an α-vector set $\bar{\Omega}_{a_t}$, which represents $v_t(\cdot \mid a_t)$, is constructed as follows:

$$\bar{\Omega}_{a_t} = (\bigoplus_{\theta_t \in \Theta} \Omega_{a_t}^{\theta_t}(x)),$$

where $\bigoplus$ is the notation of *Minkowski summation* of two sets defined as follows:

$$A \bigoplus B = \{a + b | \forall a \in A, \forall b \in B\}.$$

Again, $\bar{\Omega}_{a_t}$ is pruned to its minimal size, Ω_{a_t}, and finally, $\bar{\Omega}_t$ is constructed as:

$$\bar{\Omega}_t = \bigcup_{a_t \in A} \Omega_{a_t}.$$

$\bar{\Omega}_t$ is then pruned to a minimal size, Ω_t, to represent $v_t(\cdot)$. As discussed in Chapter 2, the purging process for an α-vector set $\bar{\Omega}_t$ involves solving $|\bar{\Omega}_t|$ linear programs. Thus, the process of

generating the minimal α-vector set is computationally challenging. In fact, the computational complexity of a general POMDP is known to be NP-hard since $|\bar{\Omega}_t| = O(|A||\Omega_{t+1}|^{|\Theta|})$ [62].

In our POMDP, if a patient has a cystoscopy and if a positive result is observed, it is assumed that treatment is triggered and the patient leaves the decision process. In other words, the minimal α-vector set $\Omega_{a_t=I}^{\theta_t=p}$ consists of just one α-vector, which corresponds to the decision to do no further cystoscopies. Therefore, $|\Omega_{a_t=I}| = O(|\Omega_{a_t=I}^{\theta_t=p}||\Omega_{a_t=I}^{\theta_t=n}|) = O(|\Omega_{a_t=I}^{\theta_t=n}|) = O(|\Omega_{t+1}|)$. If a patient defers a cystoscopy, then only a negative result is observed. Therefore, $\Omega_{a_t=W} = \Omega_{a_t=W}^{\theta_t=n}$ such that $|\Omega_{a_t=W}| = |\Omega_{a_t=W}^{\theta_t=n}| = O(|\Omega_{t+1}|)$. It follows that, $|\Omega_t| = O(|\Omega_{a_t=I}|) + O(|\Omega_{a_t=W}|) = O(|\Omega_{t+1}|)$, i.e. the size of the α- vector set for our POMDP increases polynomially with respect to the number of decision epochs.

4.4 Results

We used incremental pruning to solve our POMDP. Our implementation was developed in C++ using ILOG 12 CPLEX Concert Technology to solve the linear programs. Computational experiments were performed on a Linux server with quad core 2.83GHz CPU and 8GB RAM. In most scenarios presented below, the exact solutions were generated within 20 minutes.

4.4.1 Data Sources

The POMDP model inputs are based on the Markov model as described in Chapter 3. The base case scenario is defined in Table 4.1. We use this base case scenario to compare the optimal policy with EAU and AUA guidelines. Note that this base case is similar to that used in Chapter 3; one exception is that the disutility is set equal to 0.0003, which corresponds to the base case value used in Kulkarni [41]. We also report the results of sensitivity analysis with respect to this base case. The results, if not otherwise specified, are based on the base case scenario.

4.4.2 Base Case Scenario Results

We report results for base case scenarios for both male and female patients respectively. We compare the optimal policy with the EAU guideline and the AUA guideline for both male and female low risk patients aged 73 in the base case, as shown in Table 4.2. We selected age 73 because it is the median age for diagnosis of bladder cancer in the U.S. [1]. The EAU guideline results in the same expected QALYs compared to the more intense AUA guideline for low risk BC patients. In other words, the EAU guideline dominates the AUA guideline in that it results in the same QALYs and fewer cystoscopies.

Table 4.1: The base case model parameters for evaluation and comparison of bladder cancer surveillance strategies.

Parameter	Value	Source
μ_C	0.003	Kulkarni (2007) [41]
μ_T	0.1	
μ_{BCG}	0.08	
r_{LRDF}	0.95	
r_{IRBC}	0.95	
r_{IRDF}	0.95	
r_{HRBC}	0.95	
r_{HRDF}	0.95	
r_{MIBC}	0.80	
f	0.95	Grossman (2007) [22]
g	1	

The top part of Table 4.2 shows that the optimal policy results in a 0.13 QALY gain over the EAU and AUA guidelines for a 73 year old male patient; it also shows that the optimal policy is more intensive than the EAU guideline but less intensive than the AUA guideline. The lower part of Table 4.2 indicates that the optimal policy results in a 0.13 QALY gain over the EAU and AUA guidelines for a 73 year old female patient; it also indicates that intensity of the optimal policy, in terms of number of cystoscopies, is between the EAU and AUA guidelines. For both 73 year old male and female patients, the optimal policies resulted in 0.16 life year gain compared with the AUA and EAU guidelines.

4.4.3 Sensitivity Analysis

We performed one-way sensitivity analysis on several of the model parameters including disutility of cystoscopy, BC mortality, other cause mortality, sensitivity of cystoscopy, disutility of each treatment type, and utilities after treatment for 73 year old male patients. Sensitivity analysis is performed with respect to the base case scenario. Specifically, we consider the following scenarios:

- *Disutility of Cystoscopy*: with other parameters fixed, we changed the disutility of cystoscopy from a base case of 0.003 to a lower bound of 0.0015 (50% of the base case value), and an upper bound of 0.05 respectively. The base case value and the upper bound on disutility of cystoscopy were chosen from Kulkarni's study [41], because it is the only published study we have found that provides plausible estimates of the disutility of cystoscopy; although the estimate was based on data for cardiac catheterization procedures.

Table 4.2: The optimal policy is compared to the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected life years (95% CI), and the number of cystoscopies (95% CI). CI=Confidence Interval.

Base case; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	11.16 (11.12, 11.20)	11.91 (11.86, 11.96)	14.91 (14.86, 14.97)
EAU	11.16 (11.12, 11.20)	11.89 (11.84, 11.94)	8.40 (8.35, 8.45)
Optimal	**11.29 (11.25, 11.34)**	**12.07 (12.02, 12.11)**	**8.74 (8.67, 8.81)**

Base case; Female, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	13.03 (12.99, 13.08)	13.91 (13.86, 13.96)	15.67 (15.61, 15.73)
EAU	13.03 (12.98, 13.08)	13.88 (13.83, 13.93)	8.93 (8.87, 8.98)
Optimal	**13.16 (13.11, 13.21)**	**14.07 (14.02, 14.13)**	**9.98 (9.89, 10.06)**

- *BC Mortality*: With other parameters fixed, we changed the monthly BC mortality rate from a base case of 0.01083 (5 year DSS survival = 0.52) drawn from the study of Madersbacher (2003) [48] to a lower bound 0.005 (50% of the base case value, 5 year DSS survival = 0.74) and an upper bound of 0.01643 (150% of the base case value, 5 year DSS survival = 0.37).

- *Other Cause Mortality*: With other parameters fixed, we changed the other cause mortality from a base case drawn from CDC (2009) [26] to a lower bound of base case (50% of the base case value), and an upper bound of base case (150% of the base case value), respectively.

- *Sensitivity of Cystoscopy*: With other parameters fixed, we changed the sensitivity of cystoscopy from a base case value of 95% drawn from the study of Grossman (2007) [22] to a lower bound of 90%, and an upper bound of 100%, respectively.

- *Disutility of TURBT*: With other parameters fixed, we changed the disutility of TURBT from a base case value of 0.10 drawn from Kulkarni's study [41] to a lower bound of 0.50 (50% of the base case value), and an upper bound of 0.15 (150% of the base case value), respectively.

- *Disutility of BCG*: With other parameters fixed, we changed the disutility of BCG treatment from a base vase value of 0.08 drawn from Kulkarni's study [41] to a lower bound of 0.04 (50% of the base case value), and an upper bound of 0.12 (150% of the base case value), respectively.

- *Utility after Treatment of NMIBC*: With other parameters fixed, we changed the utility after treatment of NMIBC disease from a base case value of 0.95 drawn from Kulkarni's study [41] to a lower bound of 0.90, and an upper bound of 1, respectively.

- *Utility after Treatment of MIBC*: With other parameters fixed, we changed the utility after treatment of MIBC disease from a base case value of 0.80 drawn from Kulkarni's study [41] to a lower bound of 0.75 (50% of the base case value), and an upper bound of 0.85, respectively.

The parameter ranges in each of the defined scenarios are presented in Table 4.3.

Table 4.3: The base case values and ranges for the parameters that changed in the one-way sensitivity analysis. (Note that parameter $\delta(t)$ is dependent on age.)

Description	**Parameter**	**Base Case Value**	**Lower Bound**	**Upper Bound**
Disutility of Cystoscopy	μ_C	0.003	0.0015	0.05
Bladder Cancer Mortality	δ_{BC}	0.011	0.005	0.016
Other Cause Mortality	$\delta(t)$	$\delta(t)$	$0.5 \times \delta(t)$	$1.5 \times \delta(t)$
Sensitivity of Cystoscopy	f	0.95	0.90	1.00
Disutility of TURBT	μ_T	0.1	0.05	0.15
Disutility of BCG	μ_{BCG}	0.08	0.04	0.12
Utility after Treatment of NMIBC	r_{LRDF}	0.95	0.90	1.0
	r_{IRBC}	0.95	0.90	1.0
	r_{IRDF}	0.95	0.90	1.0
	r_{HRBC}	0.95	0.90	1.0
	r_{HRDF}	0.95	0.90	1.0
Utility after Treatment of MIBC	r_{MIBC}	0.80	0.75	0.85

The results of our sensitivity analysis are presented in Table 4.4 to Table 4.11. We observed that expected QALYs for the optimal surveillance strategy is most sensitive to other cause mortality. We also observed that the number of cystoscopies is highly sensitive to the disutility of cystoscopy and other cause mortality. A detailed discussion of the sensitivity analysis for each parameter is presented below.

Table 4.4 presents the comparison of the optimal policy with the EAU and EAU guidelines for the lower bound and the upper bound scenarios with respect to disutility of cystoscopy. The frequency of cystoscopies in the optimal policy is highly sensitive to the disutility of cystoscopy. Varying the disutility of cystoscopy from 0.05 to 0.0015 resulted in changing the number of cystoscopies from 0 to 11.20. In the lower bound scenarios, the optimal policy resulted in a 0.12 QALY gain and a 0.13 life year gain over the EAU and AUA guidelines. In the upper bound

scenario, the optimal policy resulted in a 0.37 QALY gain over the EAU and AUA guidelines; however it resulted in reduction of 0.10 life year over these two guidelines.

Table 4.4: Sensitivity analysis with respect to disutility of cystoscopy comparing the optimal policy with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected life years (95% CI), and the number of cystoscopies (95% CI). CI=Confidence Interval.

Disutility of cystoscopy, upper bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	10.46 (10.42 , 10.50)	11.91 (11.86 , 11.96)	14.91 (14.86 , 14.97)
EAU	10.77 (10.73 , 10.81)	11.89 (11.84 , 11.94)	8.40 (8.35 , 8.45)
Optimal	**11.14 (11.09 , 11.18)**	**11.79 (11.74 , 11.84)**	**0.00 (0.00 , 0.00)**

Disutility of cystoscopy, lower bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	11.18 (11.14, 11.23)	11.91 (11.86, 11.96)	14.91 (14.86, 14.97)
EAU	11.18 (11.13, 11.22)	11.89 (11.84, 11.94)	8.40 (8.35, 8.45)
Optimal	**11.30 (11.26, 11.35)**	**12.07 (12.03, 12.12)**	**11.20 (11.12, 11.27)**

Table 4.5 presents the sensitivity analysis with respect to bladder cancer mortality. Variation of the base case value of bladder cancer mortality by 50% resulted in changing the number of cystoscopies in the optimal policy from 8.69 to 6.69. In the upper bound scenario, the optimal policy resulted in a 0.13 QALY gain and 0.16 life year gain over the EAU and AUA guidelines; while in the lower bound scenario, the optimal policy results in a 0.08 QALY gain and 0.09 life year gain over the EAU and AUA guidelines. The results indicated that the worse the bladder cancer mortality the higher the incremental benefit of the optimal surveillance policy compared with the EAU and AUA guidelines.

Table 4.6 shows the sensitivity analysis with respect to other cause mortality. The results show that the frequency of cystoscopies in the optimal policy is highly sensitive to other cause mortality. Varying the base case value of other cause mortality from 150% of the base case to 50% of the base case would result in changing the mean number of cystoscopies in the optimal policy from 5.20 to 14.00. In the upper bound scenario, the optimal policy resulted in 0.01 QALY gain and 0.02 life year gain over the EAU and AUA guidelines; while in the lower bound scenario, the optimal policy results in a 0.35 QALY gain and 0.33 life year gain over the EAU and AUA guidelines. The results suggest that, for patients with lower other cause mortality, the optimal surveillance policy results in higher incremental benefits (in both QALY gain and life year gain) compared with the EAU and AUA guidelines. The results also indicated that

Table 4.5: Sensitivity analysis with respect to bladder cancer mortality comparing the optimal policy with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected life years (95% CI), and the number of cystoscopies (95% CI). CI=Confidence Interval.

Bladder cancer mortality, upper bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	11.13 (11.08, 11.17)	11.87 (11.82, 11.92)	14.92 (14.86, 14.97)
EAU	11.13 (11.09, 11.17)	11.85 (11.80, 11.89)	8.40 (8.35, 8.45)
Optimal	**11.26 (11.22, 11.31)**	**12.03 (11.98, 12.08)**	**8.69 (8.62, 8.76)**

Bladder cancer mortality, lower bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	11.22 (11.17, 11.26)	11.98 (11.94, 12.03)	14.91 (14.86, 14.97)
EAU	11.23 (11.19, 11.27)	11.97 (11.92, 12.02)	8.41 (8.36, 8.46)
Optimal	**11.31 (11.26, 11.35)**	**12.07 (12.03, 12.12)**	**6.69 (6.62, 6.76)**

patients with higher other cause mortality should generally follow less intensive surveillance.

Table 4.7 presents the sensitivity analysis with respect to sensitivity of cystoscopy. Assuming cystoscopy has a sensitivity of 100%, i.e. perfect sensitivity, the optimal policy would result in a 0.12 QALY gain and a 0.14 life year gain over the EAU and AUA guidelines by performing 7.78 cystoscopies on average; Assuming cystoscopy has a sensitivity of 90%, on the other hand, the optimal policy would result in a 0.13 QALY gain and a 0.16 life year gain over the EAU and AUA guidelines by performing 8.74 cystoscopies on average. The results suggest the optimal policy is relatively insensitive to the sensitivity of cystoscopy and the optimal policy would result in similar incremental benefits over the EAU and AUA guidelines with fewer number of cystoscopies for sensitivity in the range of 90% to 100%.

Table 4.8 presents the sensitivity analysis with respect to the disutility of TURBT. In the upper bound scenario, the optimal policy resulted in a 0.10 QALY gain and a 0.13 life year gain over the EAU and AUA guidelines with 7.80 cystoscopies on average; while in the lower bound scenario, the optimal policy would result in a 0.15 QALY gain and 0.16 life year gain over the EAU and AUA guidelines with 8.74 cystoscopies on average. The results suggest that the lower the disutility of TURBT, the more incremental benefit the optimal surveillance policy can add compared with the EAU and AUA guidelines. This is intuitive because the benefit of surveillance is ultimately influenced by the expected rewards for detection and treatment.

Table 4.9 presents the sensitivity analysis with respect to disutility of TURBT. The results are quite similar to the sensitivity analysis with respect to disutility of BCG, listed in Table 4.8. The results indicated that the lower the disutility of BCG, the more incremental benefit

Table 4.6: Sensitivity analysis with respect to other cause mortality comparing the optimal policy with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected life years (95% CI), and the number of cystoscopies (95% CI). CI=Confidence Interval.

Other cause mortality, upper bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	8.72 (8.69, 8.76)	9.31 (9.27, 9.34)	13.83 (13.78, 13.87)
EAU	8.74 (8.70, 8.77)	9.30 (9.26, 9.33)	7.63 (7.59, 7.67)
Optimal	**8.76 (8.72, 8.79)**	**9.33 (9.29, 9.36)**	**5.20 (5.15, 5.25)**

Other cause mortality, lower bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	16.88 (16.81, 16.95)	18.04 (17.97, 18.11)	16.40 (16.32, 16.47)
EAU	16.86 (16.79, 16.93)	17.99 (17.92, 18.07)	9.56 (9.50, 9.63)
Optimal	**17.23 (17.16, 17.30)**	**18.47 (18.39, 18.55)**	**14.00 (13.88, 14.12)**

Table 4.7: Sensitivity analysis with respect to sensitivity of cystoscopy comparing the optimal policy with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected life years (95% CI), and the number of cystoscopies (95% CI). CI=Confidence Interval.

Sensitivity of cystoscopy, upper bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	11.16 (11.12, 11.20)	11.91 (11.86, 11.96)	14.91 (14.86, 14.97)
EAU	11.16 (11.12, 11.21)	11.89 (11.84, 11.94)	8.40 (8.35, 8.45)
Optimal	**11.28 (11.24, 11.33)**	**12.05 (12.00, 12.10)**	**7.78 (7.71, 7.86)**

Sensitivity of cystoscopy, lower bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	11.16 (11.12, 11.20)	11.91 (11.86, 11.96)	14.91 (14.86, 14.97)
EAU	11.16 (11.12, 11.21)	11.89 (11.84, 11.94)	8.40 (8.35, 8.45)
Optimal	**11.29 (11.25, 11.34)**	**12.07 (12.02, 12.12)**	**8.74 (8.67, 8.82)**

Table 4.8: Sensitivity analysis with respect to disutility of TURBT comparing the optimal policy with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected life years (95% CI), and the number of cystoscopies (95% CI). CI=Confidence Interval.

Disutility of TURBT, upper bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	11.14 (11.09, 11.18)	11.91 (11.86, 11.96)	14.91 (14.86, 14.97)
EAU	11.14 (11.10, 11.19)	11.89 (11.84, 11.94)	8.40 (8.35, 8.45)
Optimal	**11.24 (11.20, 11.29)**	**12.04 (12.00, 12.09)**	**7.80 (7.73, 7.87)**

Disutility of TURBT, lower bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	11.18 (11.14, 11.23)	11.91 (11.86, 11.96)	14.91 (14.86, 14.97)
EAU	11.18 (11.14, 11.23)	11.89 (11.84, 11.94)	8.40 (8.35, 8.45)
Optimal	**11.33 (11.28, 11.37)**	**12.07 (12.02, 12.11)**	**8.74 (8.67, 8.81)**

the optimal surveillance policy can add compared with the EAU and AUA guidelines. Similar to the results for disutility of TURBT this is also intuitive.

Tables 4.10 and 4.11 present the sensitivity analysis with respect to the utility after treatment of NMIBC disease and the utility after treatment of MIBC disease, respectively. Changing the utility after treatment of NMIBC disease from 1 to 0.9 would result in changing the expected number of cystoscopies from 8.74 to 7.80; it also results in a drop of incremental QALYs gained from 0.14 to 0.10, and a drop of life years gained from 0.16 to 0.13, compared with the EAU and AUA guidelines. Changing the utility after treatment of MIBC disease from 0.85 to 0.75 would result in changing the number of cystoscopies from 7.80 to 8.74; it would result in an increase of incremental QALYs gained from 0.10 to 0.14 and the increase of life years gained from 0.13 to 0.16, compared with the EAU and AUA guidelines. The comparison between the sensitivity analysis with respect to utility after treatment of NMIBC disease and that after treatment of MIBC disease shows that the larger the utility loss when NMIBC disease progresses to MIBC disease, the higher the incremental benefit (in both QALY gain and life year gain) of the optimal policy over the EAU and AUA guidelines.

In summary, the optimal surveillance policy is highly sensitive to the disutility of cystoscopy and other cause mortality. Specifically, varying the disutility of cystoscopy from 150% of the base case value to 50% of the base case value would result in an increase in the expected number of cystoscopies from 7.80 to 11.20. We also found that the optimal surveillance policy can be affected by other cause mortality significantly. Changing other cause mortality from 150% of the base case to 50% of the base case would result in an increase of QALYs gained from 0.02 to

Table 4.9: Sensitivity analysis with respect to disutility of BCG comparing the optimal policy with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected life years (95% CI), and the number of cystoscopies (95% CI). CI=Confidence Interval.

Disutility of BCG, upper bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	11.14 (11.10, 11.19)	11.91 (11.86, 11.96)	14.91 (14.86, 14.97)
EAU	11.15 (11.10, 11.19)	11.89 (11.84, 11.94)	8.40 (8.35, 8.45)
Optimal	**11.25 (11.21, 11.29)**	**12.04 (12.00, 12.09)**	**7.80 (7.73, 7.87)**

Disutility of BCG, lower bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	11.18 (11.13, 11.22)	11.91 (11.86, 11.96)	14.91 (14.86, 14.97)
EAU	11.18 (11.14, 11.22)	11.89 (11.84, 11.94)	8.40 (8.35, 8.45)
Optimal	**11.32 (11.28, 11.37)**	**12.07 (12.02, 12.11)**	**8.74 (8.67, 8.81)**

Table 4.10: Sensitivity analysis with respect to utility after treatment of NMIBC comparing the optimal policy with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected life years (95% CI), and the number of cystoscopies (95% CI). CI=Confidence Interval.

Utility after treatment of NMIBC, upper bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	11.74 (11.70, 11.79)	11.91 (11.86, 11.96)	14.91 (14.86, 14.97)
EAU	11.75 (11.70, 11.79)	11.89 (11.84, 11.94)	8.40 (8.35, 8.45)
Optimal	**11.89 (11.84, 11.94)**	**12.07 (12.02, 12.11)**	**8.74 (8.67, 8.81)**

Utility after treatment of NMIBC, lower bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	10.57 (10.53, 10.62)	11.91 (11.86, 11.96)	14.91 (14.86, 14.97)
EAU	10.58 (10.54, 10.62)	11.89 (11.84, 11.94)	8.40 (8.35, 8.45)
Optimal	**10.68 (10.64, 10.72)**	**12.04 (12.00, 12.09)**	**7.80 (7.73, 7.87)**

Table 4.11: Sensitivity analysis with respect to utility after treatment of MIBC comparing the optimal policy with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected life years (95% CI), and the number of cystoscopies (95% CI). CI=Confidence Interval.

Utility after treatment of MIBC, upper bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	11.17 (11.13, 11.21)	11.91 (11.86, 11.96)	14.91 (14.86, 14.97)
EAU	11.18 (11.13, 11.22)	11.89 (11.84, 11.94)	8.40 (8.35, 8.45)
Optimal	**11.28 (11.24, 11.33)**	**12.04 (12.00, 12.09)**	**7.80 (7.73, 7.87)**

Utility after treatment of MIBC, lower bound; Male, aged 73			
Policy	QALYs	Life Years	Number of Cystoscopies
AUA	11.15 (11.10, 11.19)	11.91 (11.86, 11.96)	14.91 (14.86, 14.97)
EAU	11.15 (11.11, 11.20)	11.89 (11.84, 11.94)	8.40 (8.35, 8.45)
Optimal	**11.29 (11.24, 11.33)**	**12.07 (12.02, 12.11)**	**8.74 (8.67, 8.81)**

0.35 and an increase of life years gained from 0.02 to 0.43; it also resulted in an increase in the number of cystoscopies from 5.20 to 14.00 on average. Therefore, we conclude that patients with higher than typical other cause mortality should generally follow less intensive surveillance.

4.5 Discussion

Based on the results, we observed that the EAU guideline dominates the AUA guideline for the male and female base cases. Specifically, the EAU guideline resulted in very similar QALYs and many fewer cystoscopies compared with the AUA guideline. Furthermore, the optimal policy can result in significant QALY gain and life year gain compared with both the EAU and AUA guidelines for both male and female patients in the base case. For example, for a 73 year old male patient, the optimal policy resulted in a 0.13 QALY gain and 0.16 life year gain over the EAU and AUA guidelines; for a 73 year old female patient, the optimal policy also resulted in a 0.13 QALY gain and 0.16 life year gain over the EAU and AUA guidelines. We also observed that the optimal surveillance policy for male patients is less intensive than that for female patients of the same age. Specifically, in the base case, 73 year old male patients have 8.74 cystoscopies versus 9.98 cystoscopies on average. The only difference between male and female patients in our model is that male patients have higher other cause mortality risk. Thus, our findings suggest that older patients or patients with other competing risks should generally have less intensive surveillance.

Based on sensitivity analysis, we found that the optimal frequency of cystoscopy is highly

sensitive to the disutility of cystoscopy. Patients should undergo more intensive surveillance when the disutility of cystoscopy is reduced. For example, changing the disutility of cystoscopy from 0.05 to 0.0015 would result in an increase in the mean number of cystoscopies from 0 to 11.20. Therefore, the optimal surveillance strategy should be influenced by an individual patient's perception of the disutility of cystoscopy. Finally, changes in other cause mortality also affect the optimal policy. We observed that patients have more cystoscopies when other cause mortality drops. Specifically, changing other cause mortality from 150% to 50% of the base case would result in an increase in the mean number of cystoscopies from 5.20 to 14.00 for a 73 year old male patient. This implies that patients having comorbidity should undergo less intensive surveillance.

4.6 Conclusions

Current published guidelines, such as the EAU and AUA, are not consistent about the frequency of cystoscopic surveillance for low risk bladder cancer patients. We used a POMDP model to investigate the optimal surveillance policy that maximizes a patient's expected QALYs. Our numerical results for 73 year old male and female base case patients show that the optimal policies can result in a significant gain in QALYs (0.13 QALY gain) compared with the EAU and AUA guidelines.

Current guidelines, such as EAU and AUA, do not distinguish patients on the basis of age, comorbidity, or patient's disutility associated with cystoscopy. We performed sensitivity analysis to determine the model parameters that most affect the optimal policy. We observed that the frequency of the optimal surveillance policy is highly sensitive to the disutility of cystoscopy. For example, changing the disutility of cystoscopy from 0.05 (1,667% of the base case value) to 0.0015 (50% of the base case value) resulted in an increase in the number of cystoscopies from 0 to 11.20 on average. This suggests that a patient's preference for cystoscopy should be considered in designing the optimal surveillance policy.

We found that the optimal surveillance policy can be significantly affected by other cause mortality in terms of incremental QALY gain and incremental life year gain over the EAU and AUA guidelines as well as the number of cystoscopies. Changing other cause mortality from 150% to 50% of the base case value resulted in an increase of QALYs (life years) gained from 0.02 (0.02) to 0.35 (0.43) and an increase in the mean number of cystoscopies from 5.20 to 14.00 for a 73 year old male patient. This suggests that patients with lower other cause mortality may benefit more from the optimal policy and these patients should generally follow more intensive surveillance than the average patients. Thus, for example, patients with comorbidity that influences their other cause mortality should be screened less intensively.

Chapter 5

Optimal Surveillance Protocols Involving Urine Based Biomarkers

5.1 Introduction

In Chapter 3 a Markov model for bladder cancer surveillance for the low risk patients was presented. Various surveillance policies were evaluated using Monte Carlo simulation based on the Markov model. Chapter 4 extended this Markov model to a POMDP model to study optimal surveillance policies. Results showed that optimal policies were dependent on a patient's gender, age, other causes of mortality, and the disutility of cystoscopy. The results also demonstrated the potential to improve upon current guidelines such as the EAU and AUA guidelines. In this chapter, we explore further opportunities to improve on the guidelines by expanding the breadth of diagnostic tests to include recently developed biomarkers.

The POMDP in Chapter 4 considered a surveillance protocol based on cystoscopy alone. Recent development of urine based biomarkers has opened up research questions about the possible use of these new diagnostic tests for surveillance of low risk bladder cancer patients. Van Rhijin, et al. (2005) [67] provides a review of urine based biomarker tests that are approved or under development. All of these biomarkers have lower sensitivity and specificity than cystoscopy. Therefore the role of cystoscopy is unlikely to be replaced with these biomarkers. Nevertheless, it is possible that biomarkers may be useful for directing the frequency of cystoscopies in the surveillance process, and could possibly reduce the impact of surveillance on patients, as proposed by Fritsche, et al [20].

In this chapter we extend our POMDP model to design two surveillance protocols to evaluate the incremental benefit of urine based biomarkers. The first protocol uses cystoscopy alone, as in Chapter 4; the second protocol uses a biomarker for the initial test, with a positive result qualifying the patient for a cystoscopy. We use the new POMDP model to compare the optimal

surveillance policy resulting from using the two different protocols. As in Chapter 4, we assume that a patient will be treated immediately following a positive cystoscopy result; after treatment patients are assumed to follow a standard surveillance guideline for intermediate and high risk patients.

We begin by formulating the POMDP based on the previous POMDP formulation in Chapter 4. Then we discuss the computational complexity and how we deal with the computational challenges of this expanded POMDP model. Next, we use the incremental pruning algorithm [8] to find the optimal surveillance policy over the course of a patient's lifetime. Using a base case, we compare the optimal policies, with and without a urine based biomarker, to analyze the incremental benefit of using the biomarker. Furthermore, we simulate the optimal policies using the simulation model in Chapter 3 to investigate easy-to-implement heuristic surveillance policies. We evaluate the heuristic policies' potential to improve upon the published EAU and AUA guidelines.

We conduct sensitivity analysis with respect to model input parameters such as the disutility of cystoscopy, the sensitivity and specificity of biomarkers, and the sensitivity of cystoscopy to determine which parameters most influence the incremental benefit of biomarker tests. We use the results of our numerical experiments to draw some insights about if and how to use biomarkers to direct the frequency of cystoscopy and the potential for improving the existing guidelines.

The remainder of this chapter is organized as follows. In Section 5.2, we describe the POMDP model formulation and structure. Next in Section 5.3 we discuss the computational complexity and illustrate heuristic methods to develop easy-to-implement surveillance policies by simulating the optimal solution from our POMDP. In Section 5.4 we present the numerical results, and in Section 5.5 we discuss the implication of our results. Finally, the conclusions of our study are summarized in Section 5.6.

5.2 POMDP Model Formulation

We illustrate the surveillance decision process in Figure 5.1. At each decision epoch, the patient decides whether to initiate a diagnostic test or defer the decision until the start of the next epoch. We first consider using cystoscopy alone for surveillance. If a cystoscopy is initiated and a positive result is observed, then treatment is triggered since a positive cystoscopic result means true positive (based on perfect specificity of cystoscopy). Next, we consider using the biomarker based surveillance protocol. If a biomarker test is initiated and a positive result is observed, then a cystoscopy will be triggered; a positive triggered cystoscopic result will automatically trigger treatment. After treatment, the patient follows a standard surveillance guideline and no more surveillance decisions are made, as assumed in Chapter 4. Otherwise, if

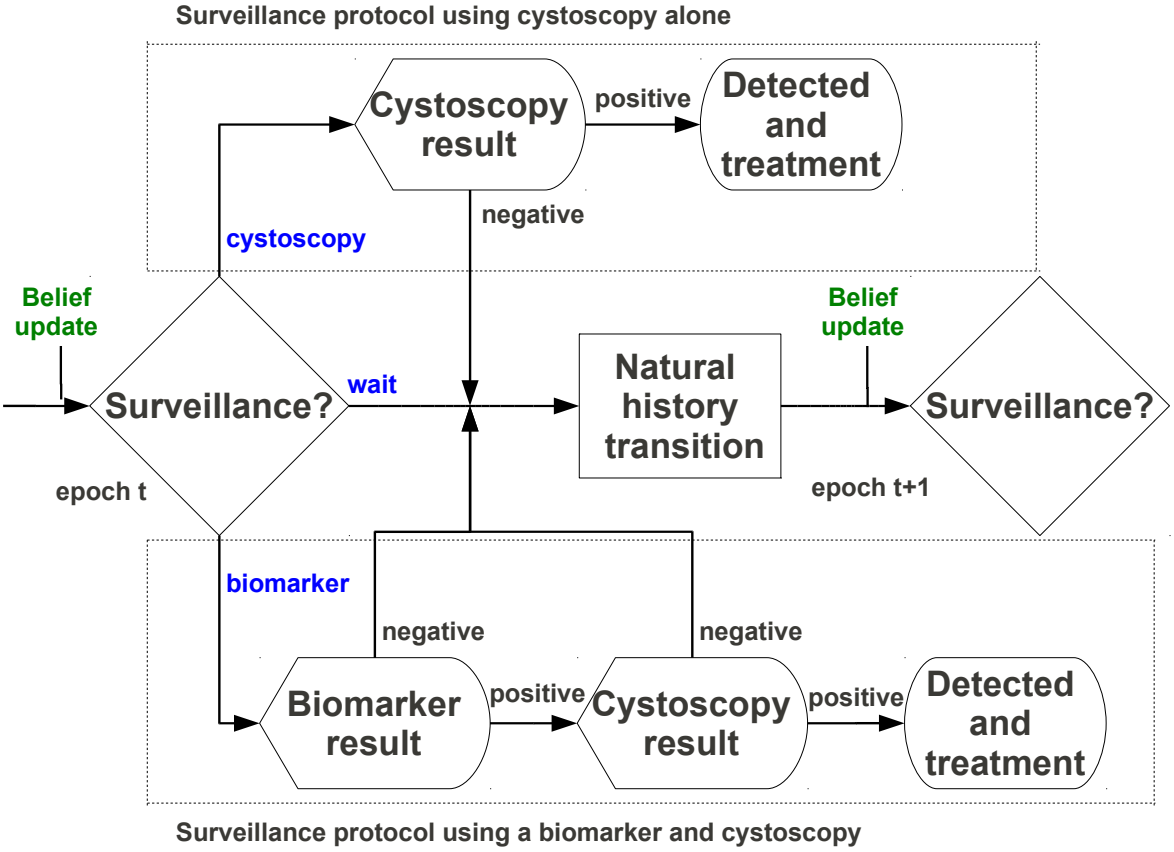

Figure 5.1: Recurring surveillance decision process for low risk bladder cancer patients with two surveillance protocols considered. The belief state is updated before choosing an action at each decision epoch.

the initial cystoscopic result is negative, or the initial biomarker test result is negative, or the triggered cystoscopic result is negative, then the patient will continue staying in the decision process and the patient's heath state evolves according to the core state transition process. If the decision is deferred, then the patient's health state also changes according to the core state transition process.

As in Chapter 4, the objective in our new POMDP model is to maximize expected QALYs. Biomarkers are assumed to have no disutility, because the procedure is non-invasive (a simple urine test conducted in the outpatient environment). However performing a biomarker test may ultimately decrease a patient's QALYs, as it may produce a positive result and trigger a cystoscopy.

We describe the POMDP model with similar notation to that used in Chapter 4. Following is a mathematical description of our model:
Stages: Surveillance is managed from a starting age (representing an initial occurrence of low risk bladder cancer) to a reasonable upper bound on life span (age 100). Decision epochs are indexed by $t = 0, 1, 2, ..., T$.

Actions: Action, $a_t \in A_t$ denotes the decision chosen at epoch t. Possible actions at each epoch t including initiating a biomarker test (B), directly initiating a cystoscopy (without first initiating a biomarker test) (C) or deferring the decision (waiting) until the next decision epoch (W). The action set is $A_t \equiv \{C, W\}, \forall t = 1, ..., T$, when we consider using cystoscopy alone. It is set as $A_t \equiv \{B, W\}, \forall t = 1, ..., T$, when we consider use the biomarker based surveillance protocol.

States: At each decision epoch a patient is in one of several health states including low risk disease free following treatment (LRDF), intermediate risk bladder cancer (IRBC), intermediate risk disease free following treatment (IRDF), high risk bladder cancer (HRBC), high risk disease free following treatment (HREF), muscle invasive bladder cancer (MIBC), and death from bladder cancer and other causes (D). We index the states by s_t and $s_t \in S = \{LRDF, IRBC, IRDF, HRBC, HRDF, MIBC, D\}$. The state description is summarized in Table 3.2. As in Chapter 4, we assume patients are treated immediately and we represent these treated states, IRDF, HRDF and MIBC, by absorbing states in our model.

Observations: After a cystoscopy ($a_t = C$) the test result is either negative (n) or positive (p). After a biomarker test ($a_t = B$) the test is either negative or positive, which automatically triggers a cystoscopy. If the triggered cystoscopy is positive, we refer to the compound observation as *double-positive* (pp). A double-positive observation confirms that both the initial biomarker test result and the triggered cystoscopic result are true positive, as the specificity of cystoscopy is assumed to be perfect. Otherwise, if the triggered cystoscopy is negative, we refer

to the compound observation as *positive-negative* (pn). A positive-negative observation indicates that either the initial biomarker test result is false positive, since cystoscopy has perfect specificity, or the initial biomarker result is true positive but the triggered cystoscopic result is false negative since cystoscopy does not have perfect sensitivity. In other words, possible observations following the urine based surveillance protocol can be negative, double-positive or positive-negative. We let n also denote the observation following the action $a_t = W$. We index the possible observations by θ_t where $\theta_t \in \Theta \equiv \{n, p, pp, pn\}$.

Information Matrix: Conditional probabilities relate the underlying core states to the observations. We let $q_t(\theta_t|s_t, a_t)$ denote the probability of observing θ_t conditional on being in core state s_t following action a_t at epoch t. We let $Q_t(a_t)$ denote the information matrix given action a_t, which has elements $q_t(\theta_t|s_t, a_t)$. As in Chapters 3 and 4, we assume that cystoscopy has perfect specificity and we use f to represent the sensitivity of cystoscopy, which is assumed to be consistent for both low grade and high grade tumors. We let φ_{lg} and φ_{hg} denote the sensitivity of a biomarker test for low grade (IRBC) and high grade (HRBC) tumors, respectively. We let τ denote the sensitivity and specificity of the biomarker test. The information matrices can be written as:

$$Q_t(C) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \begin{array}{cccc} n & p & pp & pn \end{array} \\ \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1-f & f & 0 & 0 \\ 1-f & f & 0 & 0 \end{pmatrix} \end{array},$$

$$Q_t(B) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \begin{array}{cccc} n & p & pp & pn \end{array} \\ \begin{pmatrix} \tau & 0 & 0 & 1-\tau \\ 1-\varphi_{lg} & 0 & f\varphi_{lg} & (1-f)\varphi_{lg} \\ 1-\varphi_{hg} & 0 & f\varphi_{hg} & (1-f)\varphi_{hg} \end{pmatrix} \end{array},$$

$$Q_t(W) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \begin{array}{cccc} n & p & pp & pn \end{array} \\ \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} \end{array}.$$

Belief States: As in Chapter 4, we define $\pi_t = (\pi_t(LRDF), \pi_t(IRBC), \pi_t(HRBC))$ as the belief state for the non-absorbing states at epoch t, and $\bar{\pi}_{t+1} = (\bar{\pi}_{t+1}(IRDF), \bar{\pi}_{t+1}(HRDF), \bar{\pi}_{t+1}(MIBC), \bar{\pi}_{t+1}(D))$ as the belief state that patient will enter one of the absorbing states at epoch $t+1$.

Transition Probabilities: After taking an action, an observation is observed. If an initial cysto-

scopic result is positive (p) or a triggered cystoscopic result is positive (pp), then treatment is triggered and the patient transfers to one of the post-treatment states, IRDF or HRDF. Otherwise, if an initial cystoscopic result is negative (n), or an initial biomarker test result is negative (n)s or a triggered cystoscopic result is negative (pn), then treatment is not triggered and the core state transition proceeds according to the natural process. Therefore, the state transition is only dependent on the observation at each decision epoch. We let $p_t(s_{t+1}|s_t, \theta_t)$ denote the core state transition probability from state s_t to s_{t+1}, given observation θ_t, at epoch t. We let $P_t(\theta_t)$ denote the transition probability matrix given observation θ_t, with elements $p_t(s_{t+1}|s_t, \theta_t)$ for $s_{t+1} \in S$. Similarly, we let $\bar{P}_t(\theta_t)$ denote the transition probability matrix given observation θ_t, with elements $\bar{p}_t(s_{t+1}|s_t, \theta_t)$ for $s_{t+1} \in \bar{S}$. Therefore, we have the following transition matrices defined as:

$$P_t(p) = P_t(pp) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \begin{array}{ccc} LRDF & IRBC & HRBC \end{array} \\ \begin{pmatrix} 1-\rho_{LR}-\delta(t) & \rho_{LR} & \\ & & \\ & & \end{pmatrix} \end{array},$$

$$\bar{P}_t(p) = \bar{P}_t(pp) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \begin{array}{cccc} IRDF & HRDF & MIBC & D \end{array} \\ \begin{pmatrix} & & & \delta(t) \\ 1-\delta(t) & & & \delta(t) \\ & 1-\delta(t) & & \delta(t) \end{pmatrix} \end{array},$$

and

$$P_t(n) = P_t(pn) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \begin{array}{ccc} LRDF & IRBC & HRBC \end{array} \\ \begin{pmatrix} 1-\rho_{LR}-\delta(t) & \rho_{LR} & \\ & 1-\rho_{IR}-\delta(t) & \rho_{IR} \\ & & 1-\rho_{HR}-\delta(t) \end{pmatrix} \end{array},$$

$$\bar{P}_t(n) = \bar{P}_t(pn) = \begin{array}{c} \\ LRDF \\ IRBC \\ HRBC \end{array} \begin{array}{c} \begin{array}{cccc} IRDF & HRDF & MIBC & D \end{array} \\ \begin{pmatrix} & & & \delta(t) \\ & & & \delta(t) \\ & & \rho_{HR} & \delta(t) \end{pmatrix} \end{array},$$

where ρ_{LR}, ρ_{IR}, ρ_{HR} are the recurrence rates of low risk, intermediate risk and high risk NMIBC disease, and $\delta(t)$ is other cause mortality at age t, as defined in Table 3.3. Note that empty spaces denote zeros.

Rewards: As in Chapter 4, we let $r_t(s_t, a_t)$ represent the core state expected immediate reward (measured in QALYs) given the patient is in state s_t after taking action a_t at time epoch t. Performing a cystoscopy imposes a disutility on a patient as it is an invasive test. After a cystoscopy is performed, if a positive result is observed, then treatment is triggered, imposing additional disutility associated with treatment. Thus, the belief state expected immediate reward with action a_t can be written in vector form as:

$$r_t(W) = \begin{matrix} LRDF \\ IRBC \\ HRBC \end{matrix} \begin{pmatrix} r_{LRDF} \\ r_{IRBC} \\ r_{HRBC} \end{pmatrix},$$

$$r_t(B) = \begin{matrix} LRDF \\ IRBC \\ HRBC \end{matrix} \begin{pmatrix} r_{LRDF} - (1-\tau)\mu_C \\ r_{IRBC} - \varphi(\mu_C - f(\mu_T - \mu_{BCG})) \\ r_{HRBC} - \varphi(\mu_C - f(\mu_T - \mu_{BCG})) \end{pmatrix},$$

$$r_t(C) = \begin{matrix} LRDF \\ IRBC \\ HRBC \end{matrix} \begin{pmatrix} r_{LRDF} - \mu_C \\ r_{IRBC} - \mu_C - f(\mu_T - \mu_{BCG}) \\ r_{HRBC} - \mu_C - f(\mu_T - \mu_{BCG}) \end{pmatrix},$$

where the utilities r_{LRDF}, r_{IRBC}, r_{IRDF}, r_{HRBC}, r_{HRDF}, r_{MIBC} and the disutilities of treatment μ_C, μ_T, and μ_{BCG} are as defined in Table 3.4. The expected reward vector for the absorbing states is calculated from our previous simulation model of Chapter 3 and denoted as follows:

$$R_t = \begin{matrix} IRDF \\ HRDF \\ MIBC \\ D \end{matrix} \begin{pmatrix} R_t(IRDF) \\ R_t(HRDF) \\ R_t(MIBC) \\ 0 \end{pmatrix}.$$

Similarly, we use $R_T(LRDF)$, $R_T(IRBC)$, and $R_T(HRBC)$, as in Chapter 4, to denote the expected reward at LRDF, IRBC, and HRBC at the final decision epoch T, respectively.

Although the definition of the action set and observation process is different, the optimal value function formulation of this POMDP is structurally the same as that in Chapter 4, which is written as:

$$v_t(\pi_t) = \max_{a_t \in A_t} \left\{ r_t(\pi_t, a_t) + \lambda \sum_{\theta_t \in \Theta} \{v_{t+1}(\pi_{t+1}) + \bar{\pi}_{t+1} R_{t+1}\} Pr(\theta_t | \pi_t, a_t) \right\}, \forall \pi_t, t = 1, ..., T-1, \tag{5.1}$$

and the value function of the terminal period, T, is defined as:

$$v_T(\pi_T) = \sum_{s_T \in S} \pi_T(s_T) R_T(s_T), \forall \pi_T. \tag{5.2}$$

The optimal action is defined as:

$$a_t^*(\pi_t) = \arg\max_{a_t \in A_t} \left\{ r_t(\pi_t, a_t) + \lambda \sum_{\theta_t \in \Theta} \{v_{t+1}(\pi_{t+1}) + \bar{\pi}_{t+1} R_{t+1}\} Pr(\theta_t | \pi_t, a_t) \right\}, \forall \pi_t, t = 1, ..., T-1, \tag{5.3}$$

where

$$Pr(\theta_t | \pi_t, a_t) = \sum_{s_t \in S} \pi_t(s_t) q_t(\theta_t | s_t, a_t).$$

Bayesian updates are also the same as that in Chapter 4, defined by the following formula:

$$\pi_{t+1}(s_{t+1}) = \frac{\sum_{s_t \in S} \pi_t(s_t) q_t(\theta_t | s_t, a_t) p_t(s_{t+1} | s_t, \theta_t)}{\sum_{s_t \in S} \pi_t(s_t) q_t(\theta_t | s_t, a_t)}, \forall s_{t+1} \in S. \tag{5.4}$$

and

$$\bar{\pi}_{t+1}(s_{t+1}) = \frac{\sum_{s_t \in S} \pi_t(s_t) q_t(\theta_t | s_t, a_t) \bar{p}_t(s_{t+1} | s_t, \theta_t)}{\sum_{s_t \in S} \pi_t(s_t) q_t(\theta_t | s_t, a_t)}, \forall s_{t+1} \in \bar{S}. \tag{5.5}$$

Equations (5.4) and (5.5) are used to update the belief state of a patient based on their prior belief state and their most recent action and most recent observation.

5.3 Methodology

This extended POMDP has a similar formulation to the original POMDP of Chapter 4; however due to the increased observation set, it is much harder to solve. As we discussed in Section 4.3, the computational complexity of a general POMDP is NP-hard; nevertheless the POMDP in Chapter 4 has a α-vector set that is polynomial in size. Unfortunately, the POMDP described above does not have this same property.

5.3.1 Analysis of POMDP Model

We apply incremental pruning algorithm to solve the POMDP described above. As discussed in Chapter 4, the key idea of incremental pruning is to represent the conditional value function $v_t(\cdot \mid a_t)$ for each action $a_t \in A_t$, and then focus on the conditional value function $v_t(\cdot \mid a_t, \theta_t)$ for each observation $\theta_t \in \Theta$ individually. For each action a_t, and each observation θ_t, the incremental pruning algorithm first generates an α-vector set $\bar{\Omega}_{a_t}^{\theta_t}$ to represent $v_t(\cdot \mid a_t, \theta_t)$ from the previous α-vector set Ω_{t+1}. Then $\bar{\Omega}_{a_t}^{\theta_t}$ is pruned to a minimal set $\Omega_{a_t}^{\theta_t}$ with cardinality less or equal to that of Ω_{t+1}. Next, an α-vector set $\bar{\Omega}_{a_t}$, for representing $v_t(\cdot \mid a_t)$, is constructed and pruned to its minimal size, Ω_{a_t}. Finally, $\bar{\Omega}_t$ is constructed and then pruned to a minimal size, Ω_t, to represent $v_t(\cdot)$. As explained in Chapter 4, the computational complexity of a general POMDP is NP-hard because $|\bar{\Omega}_t| = O(|A||\Omega_{t+1}|^{|\Theta|})$ [62].

In the POMDP presented in this chapter, we analyze the computational complexity for using each of the two proposed surveillance protocols separately. We first consider using cystoscopy alone, i.e. $A_t = \{C, W\}$. If a patient has a cystoscopy, and if a positive result is observed, it is assumed that treatment is triggered and the patient leaves the decision process. In other words, the minimal α-vector set $\Omega_{a_t=C}^{\theta_t=p}$ consists of one α-vector, which corresponds to the decision to do no further cystoscopies. Therefore, $|\Omega_{a_t=C}| = O(|\Omega_{a_t=C}^{\theta_t=p}||\Omega_{a_t=C}^{\theta_t=n}|) = O(|\Omega_{a_t=C}^{\theta_t=n}|) = O(|\Omega_{t+1}|)$. If a patient defers a cystoscopy, then only a negative result is observed. Therefore, $\Omega_{a_t=W} = \Omega_{a_t=W}^{\theta_t=n}$ such that $|\Omega_{a_t=W}| = |\Omega_{a_t=W}^{\theta_t=n}| = O(|\Omega_{t+1}|)$. It follows that $|\Omega_t| = O(|\Omega_{a_t=C}|) + O(|\Omega_{a_t=W}|) = O(|\Omega_{t+1}|)$, which means the size of α-vector set grows linearly in decision epoch t.

Next, we consider using a biomarker based protocol, i.e. $A_t = \{B, W\}$. If a patient chooses to initiate a biomarker test, then there will be three possible observations, negative (n), double-positive (pp) or positive-negative (pn). Similar to $\Omega_{a_t=C}^{\theta_t=p}$, $\Omega_{a_t=B}^{\theta_t=pp}$ consists of one α-vector, since a double-positive observation will trigger treatment such that the patient would leave the decision process. Nevertheless, if the observation is negative or positive-negative, the patient will continue the decision process. Therefore, $|\Omega_{a_t=B}| = O(|\Omega_{a_t=B}^{\theta_t=pp}||\Omega_{a_t=B}^{\theta_t=pn}||\Omega_{a_t=B}^{\theta_t=n}|) = O(|\Omega_{a_t=B}^{\theta_t=pn}||\Omega_{a_t=B}^{\theta_t=n}|)$. It follows that, $|\Omega_t| = O(|\Omega_{a_t=W}| + O(|\Omega_{a_t=B}|) = O(|\Omega_{t+1}|^2)$, i.e. the size of the α-vector set grows quadratically in the α-vector set at each decision epoch.

To compensate for the additional computational burden, we use three months as one decision epoch within the first three years and we use six months as one decision epoch after three years to reduce the number of decision epochs. This a reasonable assumption, because the published guidelines such as the AUA and EAU guidelines only consider multiples of six months as the surveillance intervals after two years.

5.3.2 Heuristic Policies

The POMDP solution is defined by a set of α-vectors for each decision epoch. Collectively the α-vectors define optimal decisions as a policy defined on the continuous belief state space. In other words, for a specific belief at a given decision epoch, the optimal action is chosen as the one associated with the α-vector that results in the largest value at that epoch, as illustrated by Equation 5.3. Therefore, the optimal decision at each epoch is only dependent on a patient's updated belief, which relies on the patient's previous belief, last action, and last observation. That means, the optimal policy cannot be translated into a surveillance schedule that is only dependent on time (equivalently age), without loss of accuracy. Nevertheless, from a practical point of view it is desirable to have a purely time dependent policy (if such a policy is near optimal) since such a policy is easy to implement. Note that the AUA and EAU guidelines both have this structure.

In this subsection we use heuristic methods to obtain near-optimal policies by simulating the optimal POMDP solution in the simulation model of Chapter 3. We analyze the two protocols (a) using cystoscopy alone and (b) using cystoscopy and a biomarker. We develop a heuristic policy, named *heuristic-cyst*, for (a) and another heuristic policy, named *heuristic-bmk*, for (b). We compare these heuristic policies with the optimal policies as well as the published EAU and AUA guidelines. As we describe below, the heuristic-cyst policy is the most common pattern resulting from simulating sample paths for the optimal policy using cystoscopy alone; while the heuristic-bmk policy is the most common pattern from simulating sample paths for the optimal policy using a biomarker based surveillance protocol. In the remainder of this subsection, we discuss the methods for generating the heuristic in detail.

We first consider protocol (a) that uses cystoscopy alone for surveillance. The heuristic-cyst policy has the structure defined in Table 5.1.

Table 5.1: The structure of a heuristic policy using cystoscopy alone is defined by a set of integer intervals (in months) between two consecutive cystoscopies.

A Heuristic Policy Using Cystoscopy Alone	
Heuristic-cyst	Initial cystoscopy after a_1 months; If negative during year t, next cystoscopy after a_t months. The policy is denoted as $\{a_1, a_2, ..., a_T\}$

Given an optimal policy for a POMDP using cystoscopy alone for surveillance over T years, the heuristic-cyst policy is represented by $\{a_1, a_2, ..., a_T\}$, with a_t representing waiting period after a cystoscopy during year t, as shown in Table 5.1. Based on a simulated set of sample

paths, we record the frequencies of a range of waiting periods during each year, t, and select the most common frequency during year t as a_t to define the heuristic-cyst policy.

Next, we consider protocol (b). The heuristic-bmk policy is defined to have the form in Table 5.2. Given an optimal policy for a POMDP using a biomarker based protocol for surveillance over T years, the heuristic-bmk policy is represented by two sets of numbers $\{b_1, b_2, ..., b_T\}$ and $\{c_1, c_2, ..., c_T\}$, with b_t and c_t representing the waiting periods after after a negative biomarker test and after a negative triggered cystoscopic test during year t, respectively, as illustrated in Table 5.2. Based on a simulated set of sample paths, we record the frequencies of a range of waiting periods, corresponding to a latest observation of negative biomarker test and a latest observation of negative triggered cystoscopic test, respectively. We select the most frequent waiting period after a negative biomarker test during year t as b_t, and the most frequent waiting period after a negative triggered cystoscopic test during year t as c_t to define the heuristic-bmk policy.

Table 5.2: The structure of a heuristic policy using both cystoscopy and a biomarker is defined by two sets of integer intervals (in decision epochs) for consecutive diagnostic tests.

A Heuristic Policy Using a Biomarker Based Surveillance Protocol	
Heuristic-bmk	Initial biomarker test is at month b_1; If biomarker is negative during year t, next biomarker after b_t months; If biomarker is positive, do a cystoscopy; If cystoscopy is negative in year t, next biomarker after c_t months. The policy is denoted as $\{b_1, b_2, ..., b_T\}$ and $\{c_1, c_2, ..., c_T\}$

5.4 Results

As in Chapter 4, we used incremental pruning to solve our POMDP. The implementation was developed in C++ using ILOG 12 CPLEX Concert Technology to solve the linear programs. We ran 100,000 simulated sample paths for the optimal policy to generate a heuristic policy for each of the protocol (a) and (b). We also ran 100,000 samples to calculate the mean and 95% confidence intervals for each policy to compare with the EAU and AUA published guidelines. Computational experiments were performed on a Linux server with quad core 2.83GHz CPU and 8GB RAM. In most scenarios presented below, the optimal POMDP solutions were generated within 120 minutes.

5.4.1 Data Sources

We used the same base case scenario defined in Chapter 4, in Table 4.1, for comparison of the optimal policy using cystoscopy alone with the one using a biomarker based surveillance protocol. We report the results of sensitivity analysis with respect to this base case scenario. Furthermore, we compare the heuristic-cyst and heuristic-bmk policies for each of scenarios evaluated in sensitivity analysis. Finally, we compare several of the FDA approved biomarkers with varying sensitivities and specificities. The results, if not otherwise specified, are based on the base case scenario.

5.4.2 Base Case Results

We report base case results for both male and female patients. We compared the optimal policy using cystoscopy alone (protocol (a)) and the policy using a biomarker (protocol (b)) as well as the EAU and AUA guidelines, as shown in Table 5.3. The top part of Table 5.3 shows that both the optimal policies for protocols (a) and (b) resulted in a mean of 0.14 and 0.13 QALY gain over the EAU and AUA guidelines. We observed that using the biomarker based protocol resulted in a mean reduction of 1.16 cystoscopies with a mean of 8.97 biomarker tests compared with the optimal policy using cystoscopy alone. The lower part of Table 5.3 shows that both the optimal policies for protocol (a) and (b) resulted in a mean 0.13 QALY gain over the EAU and AUA guidelines; while the optimal policy using the biomarker based protocol resulted in no reduction of cystoscopies and 12.51 biomarker tests performed on average compared with the optimal policy using cystoscopy alone. We also observed that the total number of expected biomarker tests is much lower for the male (8.97) compared to the female patient (12.51). The only difference between male and female patients in our model is that male patients have higher other cause mortality. Therefore, our results suggest that older patients and those with other competing risks should follow surveillance less intensively given the biomarker based protocol is adopted.

The most notable point to draw from the base case results is that there is no significant difference between protocol (a) and (b) on the basis of expected QALYs.

5.4.3 Sensitivity Analysis

We performed one-way sensitivity analysis on several of the model parameters including sensitivity of the biomarker, specificity of the biomarker, disutility of cystoscopy, and sensitivity of cystoscopy. Sensitivity analysis was performed with respect to the base case scenario, defined in Table 5.4. Specifically, we consider the following scenarios:

- *Sensitivity of biomarker*: with other parameters fixed, we changed the sensitivity of the

Table 5.3: The optimal policies with and without using a urine based biomarker are compared to the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected number of cystoscopies (95% CI) and the expected number of biomarkers (95% CI). CI=Confidence Interval. Optimal-cyst denotes the optimal policy using cystoscopy alone. Optimal-bmk denotes the optimal policy using a biomarker based surveillance protocol.

Male, Aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Optimal-cyst	11.29 (11.25, 11.34)	8.74 (8.67, 8.81)	-
Optimal-bmk	11.28 (11.24, 11.33)	7.58 (7.51, 7.65)	8.97 (8.95, 8.99)
Female, Aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	13.03 (12.99, 13.08)	15.67 (15.61, 15.73)	-
EAU	13.03 (12.98, 13.08)	8.93 (8.87, 8.98)	-
Optimal-cyst	13.16 (13.11, 13.21)	9.98 (9.89, 10.06)	-
Optimal-bmk	13.16 (13.11, 13.21)	9.98 (9.90, 10.07)	12.51 (12.48, 12.53)

Table 5.4: The base case model parameters for evaluation and comparison of bladder cancer surveillance strategies.

Parameter	**Value**	**Source**
μ_C	0.003	Kulkarni (2007) [41]
μ_T	0.1	
μ_{BCG}	0.08	
r_{LRDF}	0.95	
r_{IRBC}	0.95	
r_{IRDF}	0.95	
r_{HRBC}	0.95	
r_{HRDF}	0.95	
r_{MIBC}	0.80	
f	0.95	Grossman (2007) [22]
g	1	
φ_{lg}	0.47	Van Rhijin (2005) [67]
φ_{hg}	0.80	
τ	0.59	

biomarker from a lower bound of 80% of the base case to an upper bound of 120% of the base case.

- *Specificity of biomarker*: with other parameters fixed, we changed the specificity of the biomarker from a lower bound, 80% of the base case, to an upper bound, 120% of the base case.

- *Sensitivity of cystoscopy*: with other parameters fixed, we changed the sensitivity of cystoscopy from a lower bound of 0.90 to an upper bound of 1.00.

- *Disutility of Cystoscopy*: with other parameters fixed, we changed the disutility of cystoscopy from a lower bound of 0.0015 (50% of the base case value), to an upper bound of 0.05 (1,667% of the base case value).

The results of our sensitivity analysis are presented in Tables 5.5, 5.6, 5.7, and 5.8. We observed that using a biomarker based surveillance protocol can result in the same or very similar QALY gain over the EAU and AUA guidelines, with noticeable reduction in the average number of cystoscopies compared with using cystoscopy alone for most of the tested scenarios. We also observed that the optimal policy using a biomarker based protocol is highly sensitive to the disutility of the cystoscopy. Assuming the cystoscopy has a lower disutility, the optimal policy using a biomarker based protocol would result in higher QALYs with more biomarker tests. This finding is also intuitive because a reduction in the disutility of cystoscopy allows us to perform more biomarker tests to improve the chance of catching disease recurrence without losing more QALYs imposed by the additional cystoscopies triggered by the additional biomarker tests. We also found that using a biomarker with higher specificity, the optimal policy can result in more biomarker tests. This finding is intuitive because a biomarker with higher specificity corresponds to lower false positive rate, leading to a reduction in the possibility of triggering any unnecessary cystoscopies. In other words, performing more tests with a biomarker with higher specificity may help improve the chances of catching a recurrence without triggering more unnecessary cystoscopies. Therefore, the optimal policy using a biomarker with higher specificity can ultimately result in more biomarker tests. We discuss the sensitivity analysis for each scenario in detail as below.

Table 5.5 presents the comparison of the optimal policy with and without using a biomarker with the EAU and AUA guidelines for both the upper bound and the lower bound with respect to the sensitivity of the biomarker. Changing the sensitivity of the biomarker from 80% to 120% of the base case resulted in an increase in mean QALYs gained from 11.28 to 11.30. The change also resulted in an increase in the mean number of cystoscopies from 7.33 to 7.94 and an increase in the average number of biomarker tests from 8.71 to 9.41. However, none of the changes in the expected number of cystoscopies or the expected number of biomarker tests is

statistically significant. We conjecture this is because the impact of performing more diagnostic tests using a biomarker with higher sensitivity is two-fold. On the one hand, additional biomarker tests of higher sensitivity would improve the chance of catching possible disease recurrence, thus possibly leading to more QALYs. On the other hand, more tests could trigger additional unnecessary cystoscopies since the specificity of the biomarker remains unchanged, thus resulting in a negative impact on QALYs.

Table 5.5: Sensitivity analysis with respect to biomarker sensitivity by comparing the optimal policy with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected number of cystoscopies (95% CI) and the expected number of biomarkers (95% CI). The upper bound of the sensitivity of biomarker is 120% of the base case, and the lower bound of the sensitivity of biomarker is 80% of the base case. CI=Confidence Interval. Optimal-cyst denotes the optimal policy using cystoscopy alone. Optimal-bmk denotes the optimal policy using a biomarker based surveillance protocol.

Sensitivity of biomarker, upper bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Optimal-cyst	11.29 (11.25, 11.34)	8.74 (8.67, 8.81)	-
Optimal-bmk	11.30 (11.26, 11.34)	7.94 (7.87, 8.02)	9.41 (9.38, 9.43)
Sensitivity of biomarker, lower bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Optimal-cyst	11.29 (11.25, 11.34)	8.74 (8.67, 8.81)	-
Optimal-bmk	11.28 (11.23, 11.32)	7.33 (7.26, 7.40)	8.97 (8.94, 8.99)

Table 5.6 presents the sensitivity analysis with respect to the specificity of the biomarker. Changing the specificity of the biomarker from 80% to 120% of the base case resulted in an increase in the expected QALYs from 11.28 to 11.30 for a 73 year old male patient. The variation in the sensitivity of the biomarker did not resulted in a significant change in the number of cystoscopies (7.79 versus 7.68); however it resulted in a large increase of biomarker tests from 7.67 to 12.37. This is likely because a biomarker with a higher specificity is associated with a lower false positive rate, which would reduce unnecessary cystoscopies triggered by positive biomarker tests. Therefore, additional biomarker tests with higher specificity can help

improve the chances of catching possible disease recurrence without triggering more unnecessary cystoscopies. The results indicated that given a biomarker with higher specificity, the optimal policy uses an increased number of biomarker tests.

Table 5.6: Sensitivity analysis with respect to the specificity of biomarker by comparing the optimal policy with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected number of cystoscopies (95% CI) and the expected number of biomarkers (95% CI). The upper bound of the specificity of biomarker is 120% of the base case, and the lower bound of the specificity of biomarker is 80% of the base case. CI=Confidence Interval. Optimal-cyst denotes the optimal policy using cystoscopy alone. Optimal-bmk denotes the optimal policy using a biomarker based surveillance protocol.

Specificity of biomarker, upper bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Optimal-cyst	11.29 (11.25, 11.34)	8.74 (8.67, 8.81)	-
Optimal-bmk	11.30 (11.25, 11.34)	7.68 (7.61, 7.75)	12.37 (12.35, 12.40)
Specificity of biomarker, lower bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Optimal-cyst	11.29 (11.25, 11.34)	8.74 (8.67, 8.81)	-
Optimal-bmk	11.28 (11.23, 11.32)	7.79 (7.72, 7.86)	7.67 (7.65, 7.69)

Table 5.7 presents the sensitivity analysis with respect to the sensitivity of cystoscopy. Changing the sensitivity of cystoscopy from 90% to 100% results in a change in mean QALYs from 11.28 to 11.27, an increase in the mean number of cystoscopies from 7.64 to 7.74, and an increase in the mean number of biomarker tests from 9.36 to 9.50. The results suggest that the optimal policy is relatively insensitive to the sensitivity of cystoscopy in the range of 90% to 100%. We believe this is because the design of the biomarker based surveillance protocol uses cystoscopy as a secondary test triggered by a positive biomarker test. Using this protocol, the benefit of an increase on the sensitivity of cystoscopy is limited since it would not provide any additional benefit unless an initial biomarker test is true positive.

Table 5.8 presents the sensitivity analysis with respect to the disutility of cystoscopy. The optimal policy using a biomarker based surveillance protocol is highly sensitive to the disutility

Table 5.7: Sensitivity analysis with respect to the sensitivity of cystoscopy by comparing the optimal policy with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected number of cystoscopies (95% CI) and the expected number of biomarkers (95% CI). The upper bound of the sensitivity of cystoscopy is 100%, and the lower bound of the sensitivity of cystoscopy is 80%. CI=Confidence Interval. Optimal-cyst denotes the optimal policy using cystoscopy alone. Optimal-bmk denotes the optimal policy using a biomarker based surveillance protocol.

Sensitivity of cystoscopy, upper bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Optimal-cyst	11.28 (11.24, 11.33)	7.78 (7.71, 7.86)	-
Optimal-bmk	11.27 (11.22, 11.31)	7.74 (7.67, 7.81)	9.50 (9.47, 9.52)
Sensitivity of cystoscopy, lower bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Optimal-cyst	11.29 (11.25, 11.34)	8.74 (8.67, 8.82)	-
Optimal-bmk	11.28 (11.23, 11.32)	7.64 (7.57, 7.71)	9.36 (9.33, 9.38)

of cystoscopy. Varying the disutility of cystoscopy from 0.05 (1,667% of the base case) to 0.0015 (50% of the base case) resulted in an increase in the expected number of QALYs from 11.14 to 11.30. The variation with respect to the disutility of cystoscopy also resulted in a large increase in the average number of biomarker tests from 0 to 14.96 and a large increase in the average number of cystoscopies from 0 to 10.25. The reason for this significant change of the optimal policy is likely because a reduction in the disutility of cystoscopy results in the optimal policy performing more biomarker tests to improve the chance of catching possible disease recurrence without accumulating more QALY loss imposed by additional cystoscopies triggered by the additional biomarker tests.

Table 5.8: Sensitivity analysis with respect to the disutility of cystoscopy by comparing the optimal policy with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected number of cystoscopies (95% CI) and the expected number of biomarkers (95% CI). The upper bound of the disutility of cystoscopy is 0.05 QALY, and the lower bound of the disutility of cystoscopy is 0.0015 QALY. CI=Confidence Interval. Optimal-cyst denotes the optimal policy using cystoscopy alone. Optimal-bmk denotes the optimal policy using a biomarker based surveillance protocol.

Disutility of cystoscopy, upper bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	10.46 (10.42 , 10.50)	14.91 (14.86, 14.97)	-
EAU	10.77 (10.73 , 10.81)	8.40 (8.35, 8.45)	-
Optimal-cyst	11.14 (11.09 , 11.18)	0.00 (0.00 , 0.00)	-
Optimal-bmk	11.14 (11.09 , 11.18)	0.00 (0.00 , 0.00)	0.00 (0.00 , 0.00)
Disutility of cystoscopy, lower bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.18 (11.14, 11.23)	14.91 (14.86, 14.97)	-
EAU	11.18 (11.13, 11.22)	8.40 (8.35, 8.45)	-
Optimal-cyst	11.30 (11.26, 11.35)	11.20 (11.12, 11.27)	-
Optimal-bmk	11.30 (11.26, 11.34)	10.25 (10.18, 10.32)	14.96 (14.92, 15.00)

In summary, the optimal policy using a biomarker based surveillance protocol is highly sensitive to the disutility of cystoscopy. Specifically, varying the disutility of cystoscopy from 1,667% to 50% of the base case resulted in a significant change in the average number of biomarker tests from 0 to 14.96 and an large increase in the average number of cystoscopies from 0 to 10.25. The surveillance intensity of the optimal policy using a biomarker based protocol

is highly sensitive to the specificity of the biomarker. Specifically, varying the specificity of the biomarker from 47% (80% of the base case) to 71% (120% of the base case) would result an increase in the average number of biomarker tests from 7.67 to 12.37. However, the large increase in the number of biomarker tests resulted in a reduction in the expected number of cystoscopies from 7.79 to 7.68 because the biomarker with a higher specificity triggered less unnecessary cystoscopies due to its lower false positive rate.

5.4.4 Evaluation of Heuristic Policies

In addition to comparing the optimal policies using protocol (a) and (b) with the published EAU and AUA guidelines, we also developed a heuristic policy, heuristic-cyst, for (a) and a heuristic policy, heuristic-bmk, for (b). The heuristics were used to generate policies for the male and female base case patients. We present results comparing these two heuristic policies with the EAU and AUA guidelines in Table 5.9. Furthermore, we developed a heuristic-cyst policy and a heuristic-bmk policy for each of the scenarios evaluated in the above sensitivity analysis for optimal policies. We compared these heuristic policies with the EAU and AUA guidelines from Table 5.11 to Table 5.14. We observed that all the heuristic-cyst policies have non-decreasing surveillance intervals after the first year; while not all the heuristic-bmk policies have such a property. We observed that for all the tested scenarios, the heuristic-bmk policies result in similar or higher mean QALYs compared with the corresponding heuristic-cyst policies.

Based on the sensitivity analysis, we observed that the intervals of the heuristic-cyst policy can be significantly affected by the sensitivity of cystoscopy. The higher the sensitivity of cystoscopy the larger the waiting periods because a negative result from a cystoscopy with higher sensitivity is less likely to be a false negative. As for the heuristic-bmk policy, we observed that both the waiting period after a negative biomarker test and the waiting period after a negative triggered cystoscopic test is affected by the sensitivity and specificity of the biomarker. Assuming the biomarker has higher sensitivity, a negative biomarker test would be less likely to be a false negative, therefore the waiting period after a negative biomarker test, b_t, should increase; on the other hand a positive biomarker test would be more likely to be a true positive, therefore the waiting period after a negative triggered cystoscopy test, c_t should decrease. Assuming the biomarker has a higher specificity, a negative biomarker test would be more likely to be a true negative, therefore the waiting period after a negative biomarker test, b_t, should increase; on the other hand a positive biomarker test would be less likely to be a false positive, therefore the waiting period after a negative triggered cystoscopy test, c_t should decrease. We observed each of these trends in the sample paths generated based on the optimal policies for protocols (a) and (b). We also found that both the heuristic-cyst policy and the heuristic-bmk policy are highly sensitive to the disutility of cystoscopy. The lower the disutility

of cystoscopy the more intensive these two heuristic policies and the shorter the waiting periods, a_t, b_t, and c_t for these two heuristic policies. We discuss each of the evaluated scenario in detail as follows.

Table 5.9 represents the heuristic-cyst policy and the heuristic-bmk policy for base case male and female patients. The upper part of Table 5.9 shows that both the heuristic-cyst policy and the heuristic-bmk policy result in 0.05 QALYs gain over the EAU and AUA guidelines. The heuristic-cyst policy turns out to be less invasive than both the EAU and AUA guidelines. It schedules cystoscopy yearly in the first two years; after that the surveillance interval increases from every 18 months to 60 months during year 3 to year 7, and 7 years later the patient will stop surveillance if no recurrence is detected. The heuristic-bmk policy schedules the first biomarker test at 12 months. If a biomarker test is negative, then the patient should have another biomarker test 3 months later in year 2 to 3 or 6 months later in year 4 to 9, after which patient can stop surveillance. If a positive biomarker result is observed, then a cystoscopy should be triggered immediately. A negative cystoscopy result in year 1 and year 2 results in a delay of the next biomarker test to 12 months later; while a negative cystoscopy test in year 3 to year 7 results in a delay of the next biomarker test to 12 to 24 months later. After 7 years the patient surveillance stops if there is no recurrence (as observed for the optimal policy). The two heuristic policies are summarized in Table 5.10. Finally we observed that the heuristic-bmk policy results in a saving of 0.66 cystoscopies at the cost of 8.11 biomarker tests on average. The lower part of Table 5.9 shows that the heuristic-cyst policy results in 0.09 QALY gain over the EAU and AUA guidelines and the heuristic-bmk policy results in an incremental 0.03 QALY gain over the heuristic-cyst policy. We find that the heuristic policies for female patients is quite different from the heuristic policies for male patients of the same age.

Table 5.11 shows the relationship between the heuristic-bmk policy with the sensitivity of the biomarker. We observed that in the upper bound case (120% of the base case sensitivity), each of the waiting periods after a negative biomarker test except the first one, b_t, is basically larger than the corresponding waiting periods in the lower bound case (80% of the base case sensitivity). This finding is intuitive because a negative result from a biomarker with higher sensitivity is less likely to be a false negative so that the patient should wait for a longer time period before the next biomarker test. We also observed in the upper bound case, the waiting period after a negative triggered cystoscopic test, c_t, is shorter than the corresponding waiting periods in the lower bound case. This finding is also intuitive because the positive result from a biomarker with higher sensitivity is more likely to be a true positive, which means a negative triggered cystoscopic result would be more likely to be false negative, so that the patient should wait for shorter time before the next biomarker test. We also found that changing the sensitivity of the biomarker from 80% to 120% of the base case resulted in an increase in the expected QALY gain from 11.21 to 11.23, which is not statistical significant.

Table 5.9: Comparison of the heuristic-cyst policy and the heuristic-bmk policy to the AUA and EAU guidelines for both male and female patients aged 73 in the base case. The heuristic-cyst policy is denoted as $\{a_t\}$, illustrated in Table 5.1. The heuristic-bmk policy is denoted as $\{b_t\}$ and $\{c_t\}$, illustrated in Table 5.2.

Male, Aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Heuristic-cyst	11.21 (11.17, 11.25)	7.69 (7.62, 7.76)	-
	$\{a_t\}$={12 12 18 30 30 60 60 120 120 120 120 120}		
Heuristic-bmk	11.21 (11.17, 11.26)	7.03 (6.97, 7.10)	8.11 (8.08, 8.13)
	$\{b_t\}$={ 12 3 3 6 6 6 6 6 6 30 120 120 }		
	$\{c_t\}$={ 12 12 24 18 12 24 18 48 42 120 120 120 }		

Female, Aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	13.03 (12.99, 13.08)	15.67 (15.61, 15.73)	-
EAU	13.03 (12.98, 13.08)	8.93 (8.87, 8.98)	-
Heuristic-cyst	13.12 (13.07, 13.17)	8.94 (8.86, 9.02)	-
	$\{a_t\}$={ 12 15 21 21 30 54 54 54 120 120 120 120 }		
Heuristic-bmk	13.15 (13.10, 13.20)	10.12 (10.04, 10.20)	13.12 (13.10, 13.15)
	$\{b_t\}$={ 3 3 3 6 6 12 12 6 36 30 120 120 }		
	$\{c_t\}$={ 6 9 9 6 12 12 12 18 36 30 120 120 }		

Table 5.10: The most common patterns found by simulating the optimal surveillance policies for aged 73 typical male patients in the base case.

	Heuristic Policies for Male Aged 73, Base Case
Heuristic-bmk	Initial biomarker test at 12 months; If biomarker is negative in the first 3 (9) years, next biomarker after 3 (6) months; If biomarker is positive, do a cystoscopy; If cystoscopy is negative or false positive within 2 years, next biomarker after 12 months; If cystoscopy is negative or false positive after 2 years, next biomarker after 12-24 months. Stop surveillance after 10 years if no recurrence occurs;
Heuristic-cyst	If negative within 2 years, next cystoscopy after 12 months; If negative during year 3, next cystoscopy after 18 months; Else if negative within 5 (7) years, next cystoscopy after 30 (60) months; If negative after 8 years, next cystoscopy after 21 months. Stop surveillance after 7 years if no recurrence occurs;

Table 5.12 shows the relationship between the heuristic-bmk policy and the specificity of the biomarker. We observed that in the upper bound case (120% of the base case specificity), each of the waiting periods after a negative biomarker test except the first one, b_t, is longer than the corresponding waiting periods in the lower bound case (80% of the base case specificity). This finding is intuitive because a negative result from a biomarker with higher specificity is more likely to be a true negative, therefore the patient should wait for a longer period of time before the next biomarker test. We also observed in the upper bound case that, the waiting period after a negative triggered cystoscopic test, c_t, is shorter than the corresponding waiting periods in the lower bound case. This finding is also intuitive because the positive result from a biomarker with higher sensitivity is less likely to be a false positive, which means a negative triggered cystoscopic result would be less likely to be true negative. Thus, the patient should wait for a shorter period of time before the next biomarker test. We also found that changing the specificity of the biomarker from 80% to 120% of the base case resulted in an increase in the expected QALY gain from 11.22 to 11.23, which is not statistically significant.

Table 5.13 presents the relationship between the heuristic policies, heuristic-cyst and heuristic-bmk, and the sensitivity of cystoscopy. For the heuristic-cyst policy, we observed each of the waiting periods after a negative cystoscopic test, a_t, resulted from the upper bound case (100%) is shorter than the corresponding waiting periods in the lower bound case (90%). This finding

Table 5.11: Sensitivity analysis with respect to biomarker sensitivity by comparing the heuristics policies, heuristic-cyst and heuristic-bmk, with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected number of cystoscopies (95% CI) and the expected number of biomarkers (95% CI). CI=Confidence Interval. The heuristic-cyst policy is denoted as $\{a_t\}$, illustrated in Table 5.1. The heuristic-bmk policy is denoted as $\{b_t\}$ and $\{c_t\}$, illustrated in Table 5.2.

Sensitivity of biomarker, upper bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Heuristic-cyst	11.21 (11.17, 11.25)	7.69 (7.62, 7.76)	-
	$\{a_t\}$={12 12 18 30 30 60 60 120 120 120 120 120}		
Heuristic-bmk	11.23 (11.18, 11.27)	6.79 (6.72, 6.86)	7.63 (7.61, 7.64)
	$\{b_t\}$={ 9 6 6 12 12 18 12 18 36 120 120 120 }		
	$\{c_t\}$={ 9 9 9 12 12 18 12 18 36 120 120 120 }		

Sensitivity of biomarker, lower bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Heuristic-cyst	11.21 (11.17, 11.25)	7.69 (7.62, 7.76)	-
	$\{a_t\}$={12 12 18 30 30 60 60 120 120 120 120 120}		
Heuristic-bmk	11.21 (11.17, 11.26)	6.63 (6.56, 6.69)	7.82 (7.80, 7.84)
	$\{b_t\}$={ 12 3 3 6 6 6 6 6 6 24 6 6 }		
	$\{c_t\}$={ 12 12 18 12 30 24 60 54 42 30 120 120 }		

Table 5.12: Sensitivity analysis with respect to the specificity of the biomarker by comparing the heuristics policies, heuristic-cyst and heuristic-bmk, with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected number of cystoscopies (95% CI) and the expected number of biomarkers (95% CI). CI=Confidence Interval. The heuristic-cyst policy is denoted as $\{a_t\}$, illustrated in Table 5.1. The heuristic-bmk policy is denoted as $\{b_t\}$ and $\{c_t\}$, illustrated in Table 5.2.

Specificity of biomarker, upper bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Heuristic-cyst	11.21 (11.17, 11.25)	7.69 (7.62, 7.76)	-
	$\{a_t\}$={12 12 18 30 30 60 60 120 120 120 120 120}		
Heuristic-bmk	11.24 (11.19, 11.28)	7.67 (7.60, 7.74)	12.91 (12.88, 12.94)
	$\{b_t\}$={ 6 3 6 6 6 6 12 12 12 30 120 120 }		
	$\{c_t\}$={ 6 3 12 6 12 12 12 12 12 30 120 120 }		

Specificity of biomarker, lower bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Heuristic-cyst	11.21 (11.17, 11.25)	7.69 (7.62, 7.76)	-
	$\{a_t\}$={12 12 18 30 30 60 60 120 120 120 120 120}		
Heuristic-bmk	11.22 (11.17, 11.26)	7.66 (7.59, 7.73)	7.79 (7.77, 7.81)
	$\{b_t\}$={ 3 3 6 6 6 6 6 6 6 24 6 6 }		
	$\{c_t\}$={ 12 12 18 12 30 24 60 54 42 30 120 120 }		

is intuitive because a negative result from a cystoscopy with higher sensitivity is less likely to be a false negative so that the patient should wait for a longer time period before the next cystoscopy test. As to the heuristic-bmk policy, we found that in the lower bound case the waiting periods, b_t and c_t, are non-decreasing with the year, t, being disease free; however the waiting periods after a negative biomarker test, b_t, do not seem to have strong linear relationship with the year t. For example, $b_4 = 6, b_5 = 12, b_6 = 6, b_7 = 6, b_8 = 12, b_9 = 6$. It may be highly dependent on the time since last cystoscopy test since a negative triggered cystoscopic test, which has a perfect specificity, would confirm that patient is disease free with probability 100%, which may affect the Bayesian updating of patient's belief for a long time based on the Bayesian updating formula 5.4.

Table 5.13: Sensitivity analysis with respect to the sensitivity of cystoscopy by comparing the heuristics policies, heuristic-cyst and heuristic-bmk, with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected number of cystoscopies (95% CI) and the expected number of biomarkers (95% CI). CI=Confidence Interval. The heuristic-cyst policy is denoted as $\{a_t\}$, illustrated in Table 5.1. The heuristic-bmk policy is denoted as $\{b_t\}$ and $\{c_t\}$, illustrated in Table 5.2.

Sensitivity of cystoscopy, upper bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Heuristic-cyst	11.20 (11.16, 11.25)	6.63 (6.56, 6.69)	-
	$\{a_t\}$={ 12 18 18 36 60 60 60 120 120 120 120 120 }		
Heuristic-bmk	11.24 (11.19, 11.28)	8.55 (8.48, 8.62)	11.49 (11.46, 11.51)
	$\{b_t\}$={ 3 3 3 6 12 6 6 12 6 30 120 120 }		
	$\{c_t\}$={ 9 9 9 12 12 18 18 12 42 30 120 120 }		

Sensitivity of cystoscopy, lower bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Heuristic-cyst	11.20 (11.16, 11.25)	7.67 (7.60, 7.74)	-
	$\{a_t\}$={ 12 12 21 27 27 60 60 120 120 120 120 120 }		
Heuristic-bmk	11.20 (11.16, 11.25)	7.24 (7.17, 7.31)	8.94 (8.92, 8.96)
	$\{b_t\}$={ 6 6 6 6 6 12 12 12 36 120 120 120 }		
	$\{c_t\}$={ 6 6 12 12 18 24 24 48 36 120 120 120 }		

Table 5.14 presents the changes of the heuristic policies by varying the disutility of cystoscopy from 0.0015 (50% of the base case) to 0.05 (1,667% of the base case). We observe that the heuristic-cyst policy with the upper bound on cystoscopy disutility is much less intensive than that for the lower bound on cystoscopy disutility. We also observe that the heuristic-bmk policy is also highly sensitive to the disutility of cystoscopy. Specifically, varying the disutility from 0.0015 to 0.05 resulted in a decrease of number of biomarker tests from 16.97 to 0.52 and a decrease of cystoscopies from 11.09 to 1.02. This finding is very intuitive, because a higher disutility value will offset the benefit of possible early detection of recurrence.

Table 5.14: Sensitivity analysis with respect to the disutility of cystoscopy by comparing the heuristics policies, heuristic-cyst and heuristic-bmk, with the AUA and EAU guidelines in terms of the expected QALYs (95% CI), the expected number of cystoscopies (95% CI) and the expected number of biomarkers (95% CI). CI=Confidence Interval. The heuristic-cyst policy is denoted as $\{a_t\}$, illustrated in Table 5.1. The heuristic-bmk policy is denoted as $\{b_t\}$ and $\{c_t\}$, illustrated in Table 5.2.

Disutility of cystoscopy, upper bound; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	10.46 (10.42 , 10.50)	14.91 (14.86, 14.97)	-
EAU	10.77 (10.73 , 10.81)	8.40 (8.35, 8.45)	-
Heuristic-cyst	11.00 (10.95, 11.04)	2.19 (2.14, 2.23)	-
	$\{a_t\}$={ 120 120 120 120 120 120 120 120 120 120 120 120 }		
Heuristic-bmk	11.04 (10.99, 11.08)	1.02 (0.99, 1.06)	0.52 (0.52, 0.53)
	$\{b_t\}$={ 120 120 120 120 120 120 120 120 120 120 120 120 }		
	$\{c_t\}$={ 120 120 120 120 120 120 120 120 120 120 120 120 }		
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.18 (11.14, 11.23)	14.91 (14.86, 14.97)	-
EAU	11.18 (11.13, 11.22)	8.40 (8.35, 8.45)	-
Heuristic-cyst	11.24 (11.20, 11.29)	10.39 (10.32, 10.46)	-
	$\{a_t\}$={ 9 9 12 12 15 18 21 39 39 120 120 120 }		
Heuristic-bmk	11.25 (11.20, 11.29)	11.09 (11.01, 11.16)	16.97 (16.93, 17.01)
	$\{b_t\}$={ 3 3 3 6 6 6 6 12 6 30 6 6 }		
	$\{c_t\}$={ 6 3 6 6 6 6 6 12 6 30 120 120 }		

5.4.5 Comparison of Biomarker Tests

Other biomarker tests, in addition to NMP22, are being developed. For example, BTAstat, BTAtrak, FDP, ImmunoCyt and FISH (UroVysion) have all been approved by the FDA for diagnosis of bladder cancer recurrence [67]. These biomarkers have different sensitivities and specificities compared with NMP22, as listed in Table 5.15.

Table 5.15: Sensitivity, specificity of the urine based biomarkers that are approved by FDA for use in diagnosing bladder cancer recurrence [67].

	NMP22	BTAstat	BTA Trak	FDP	Immunocyt	FISH
Sensitivity for low and intermediate risk NMIBC	47%	56%	57%	63%%	84%	67%
Sensitivity for high risk NMIBC	80%	75%	74%	86%	100%	95%
specificity for low, intermediate and high risk NMIBC	59%	79%	66%	80%	62%	47%

We evaluated each of the FDA approved biomarkers listed above. We compared the resulting optimal surveillance policies for the base case scenario, as shown in Table 5.16. We observed that the optimal policy using each of the six biomarkers resulted in significant QALY gain (from 0.14 to 0.19) over the EAU and AUA guidelines. However, none of the biomarker tests led to statistically significant improvements in QALYs over the use of cystoscopy alone for the base case. We further evaluated the optimal policy using a perfect biomarker (100% sensitivity and 100% specificity) which resulted in a mean of 11.31 QALYs, which is not statistically larger than that for the optimal policy using cystoscopy alone (11.29).

Further observation from Table 5.16 for the base case for the male patient revealed the optimal policies resulted in large variation in the average number of biomarker tests, from 8.18 to 15.10. However, the average number of cystoscopies resulting from each of the optimal policies are similar (7.49 to 7.69). We further analyzed the heuristic-bmk policies using each of six biomarkers, named *Heuristic-NMP22*, *Heuristic-BTA-Stat*, *Heuristic-BTA-Trak*, *Heuristic-FDP*, *Heuristic-Immunocyt* and *Heuristic-FISH*. We compared all of these heuristic policies in Table 5.17. We observed that all of these heuristic policies using FDA approved biomarkers, except *Heuristic-NMP22*, resulted in an incremental QALY gain of 0.04 or more compared with the heuristic policy using cystoscopy alone. Furthermore, we found that the heuristic policy using a perfect biomarker (perfect sensitivity and perfect specificity) resulted in an incremental QALY gain of 0.08 compared with the heuristic policy using cystoscopy alone. These findings

indicate that the purely time dependent heuristic policy using the biomarker based surveillance protocol has the potential to significantly improve the purely time dependent heuristic policy using cystoscopy alone.

We also observed that the heuristic-bmk policies resulted in a large variation in the average number of biomarker tests, from 7.41 to 16.13; but similar average number of cystoscopies, from 7.03 to 7.67. We also observed that the waiting periods after a negative biomarker test, b_t, and the waiting periods after a negative triggered cystoscopic test, c_t, were generally non-decreasing over time. For example, given FISH is used for surveillance, the heuristic-FISH policy suggested a biomarker test every 6 months in the first four years, and yearly for another five years if no recurrence occurs.

From the analysis of the sensitivity and specificity of the biomarker in Tables 5.11 and 5.12, we observed that the waiting period after a negative biomarker test, b_t, is positively correlated with both the sensitivity and specificity of the biomarker; while we also learned that the waiting period after a negative triggered cystoscopic result, c_t, is negatively correlated with both the sensitivity and specificity of the biomarker. However, one-way sensitivity analysis cannot determine if b_t and c_t is most affected by the sensitivity or the specificity of the biomarker. We can answer this question by compare these heuristic polices using each of these biomarkers with different combinations of sensitivity and specificity. We excluded NMP22 for this cross comparison because it has a very low sensitivity for low grade and quite high sensitivity for high grade tumors, which makes it hard to compare with the sensitivities of the other biomarkers. Of the other five biomarkers, Immunocyt has the highest sensitivity and the lowest specificity, and heuristic-Immunocyt has the largest waiting periods, based on the observation of b_t and c_t. This finding indicates that b_t may be most sensitive to the sensitivity of the biomarker and c_t may be most sensitive of the specificity of the biomarker. We also found that BTA-Stat has the lowest sensitivity and the second highest specificity (79%, the highest being 80%) of the five biomarkers considered, and heuristic-BTA-Stat has the shortest waiting periods b_t and c_t. This finding strengthens the indication that b_t is most sensitive to the sensitivity of the biomarker and c_t is most sensitive of the specificity of the biomarker.

5.5 Discussion

Based on the base case results, we observed that the optimal policy using a biomarker based protocol result in no significant QALY gain over the protocol based on cystoscopy alone. For example, in the base case, for both a male patient and a female patient of age 73, the optimal policy using a biomarker based protocol resulted in 0.13 QALY gain over the EAU and AUA guidelines. However, the difference between the biomarker based protocols, in terms of expected QALYs, were not statistically significantly different from policies based on cystoscopy alone.

Table 5.16: Evaluating the optimal policies using different biomarkers

Male, Aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
AUA	11.16 (11.12, 11.20)	14.91 (14.86, 14.97)	-
EAU	11.16 (11.12, 11.21)	8.40 (8.35, 8.45)	-
Optimal-cyst	11.29 (11.25, 11.34)	8.74 (8.67, 8.81)	-
Optimal-NMP22	11.28 (11.24, 11.33)	7.58 (7.51, 7.65)	8.97 (8.95, 8.99)
Optimal-BTA-Stat	11.30 (11.25, 11.34)	7.54 (7.47, 7.62)	15.10 (15.06, 15.13)
Optimal-BTA-Trak	11.29 (11.25, 11.34)	7.69 (7.61, 7.76)	10.54 (10.51, 10.56)
Optimal-FDP	11.31 (11.26, 11.35)	7.49 (7.41, 7.56)	14.76 (14.72, 14.80)
Optimal-Immunocyt	11.30 (11.25, 11.34)	7.59 (7.52, 7.67)	8.18 (8.16, 8.20)
Optimal-FISH	11.30 (11.26, 11.35)	7.49 (7.42, 7.57)	10.43 (10.40, 10.45)
Optimal-perfect-bmk	**11.31 (11.27, 11.36)**	**5.05 (4.97, 5.13)**	**27.75 (27.68, 27.83)**

Table 5.17: Comparison of the heuristic-bmk policies using each of the FDA approved biomarkers, named *Heuristic-NMP22*, *Heuristic-BTA-Stat*, *Heuristic-BTA-Trak*, *Heuristic-FDP*, *Heuristic-Immunocyt* and *Heuristic-FISH* in terms of the expected QALYs (95% CI), the expected number of cystoscopies (95% CI) and the expected number of biomarkers (95% CI). CI=Confidence Interval. These heuristic-bmk policies is denoted as $\{b_t\}$ and $\{c_t\}$, illustrated in Table 5.2.

Base Case; Male, aged 73			
Policy	Expected QALYs	Number of Cystoscopies	Number of Biomarkers
Heuristic-NMP22	11.21 (11.17, 11.26)	7.03 (6.97, 7.10)	8.11 (8.08, 8.13)
	$\{b_t\}$={ 12 3 3 6 6 6 6 6 6 30 120 120 }		
	$\{c_t\}$={ 12 12 24 18 12 24 18 48 42 120 120 120 }		
Heuristic-BTA-Stat	11.26 (11.22, 11.31)	7.63 (7.56, 7.71)	16.13 (16.09, 16.16)
	$\{b_t\}$={ 3 3 6 6 6 6 6 6 12 30 120 120 }		
	$\{c_t\}$={ 3 9 6 6 6 6 6 6 12 30 120 120 }		
Heuristic-BTA-Trak	11.25 (11.21, 11.29)	7.67 (7.59, 7.74)	11.04 (11.01, 11.06)
	$\{b_t\}$={ 3 3 6 6 6 6 12 18 36 120 120 120 }		
	$\{c_t\}$={ 9 9 15 12 12 18 12 18 36 120 120 120 }		
Heuristic-FDP	11.25 (11.21, 11.30)	7.55 (7.48, 7.63)	15.55 (15.51, 15.59)
	$\{b_t\}$={ 6 3 6 6 6 6 6 6 6 30 120 120 }		
	$\{c_t\}$={ 6 3 9 6 6 6 6 6 6 30 120 120 }		
Heuristic-Immunocyt	11.25 (11.20, 11.29)	7.20 (7.13, 7.27)	7.41 (7.39, 7.43)
	$\{b_t\}$={ 9 9 9 12 12 12 12 18 36 120 120 120 }		
	$\{c_t\}$={ 9 9 9 12 12 12 12 18 36 120 120 120 }		
Heuristic-FISH	11.25 (11.20, 11.29)	7.44 (7.36, 7.51)	10.61 (10.59, 10.64)
	$\{b_t\}$={ 6 6 6 6 12 12 12 12 12 30 120 120 }		
	$\{c_t\}$={ 6 6 6 6 12 12 12 12 12 30 120 120 }		
Heuristic-perfect-bmk	**11.29 (11.25, 11.34)**	**4.90 (4.82, 4.98)**	**27.10 (27.03, 27.18)**
	$\{b_t\} = \{3, 3, 3, 3, 3, 3, 3, 3, 3, 3, 33, 120, 120\}$		
	$\{c_t\} = \{3, 120, 120, 120, 120, 120, 120, 120, 120, 120, 120, 120, 120\}$		

We also observed that the optimal policy using a biomarker based surveillance protocol for male patients is less intensive than that for female patients of the same age. Specifically, in the base case, the total number of expected biomarker tests for a 73 year old male patient is 8.97 versus 12.51 for a 73 year old female patient. In our model, the only difference between male and female patients of the same age is that male patients have higher other cause mortality. Thus it implies that older patients or patients with other competing risks should generally have less intensive surveillance given a biomarker based protocol is adopted.

Based on the sensitivity analysis, we found that the value function for the optimal policies for protocols (a) and (b) were not statistically significantly different. However, the optimal surveillance frequency using a biomarker based surveillance protocol is highly sensitive to the specificity of the biomarker. Patients should follow more intensive surveillance if a biomarker with higher specificity is used in surveillance. Specifically, varying the specificity of the biomarker from 47% (80% of the base case) to 71% (120% of the base case) resulted an increase of biomarker tests from 7.67 to 12.37 on average. We also found that the disutility of cystoscopy affected the optimal policy using a biomarker based surveillance protocol. Specifically, varying the disutility of cystoscopy from 0.0015 to 0.05 resulted in a decrease in the average number of biomarker tests from 14.96 to 0. Therefore, the optimal strategy using a biomarker based surveillance protocol should be influenced by the individual patient's perception of the disutility of cystoscopy.

Based on the sensitivity analysis, we found that the intervals of the heuristic-cyst policy are affected by the sensitivity of cystoscopy. Patient should wait for a longer time period given a negative result from a cystoscopy with higher sensitivity is observed. Based on the sensitivity analysis on heuristic-bmk we observed that the waiting period after a negative biomarker test, b_t, is positively correlated with both the sensitivity and specificity of the biomarker; we also observed that the waiting period after a negative triggered cystoscopic result, c_t, is negatively correlated with both the sensitivity and specificity of the biomarker.

We compared the optimal policy using a perfect biomarker (perfect sensitivity and perfect specificity) with the optimal policy using cystoscopy alone. We found that using the perfect biomarker (protocol (b)) did not result in statistically larger QALY gain compared to the optimal policy without using a biomarker (protocol (a)). We also compared the heuristic policy using a perfect biomarker with the heuristic policy using cystoscopy alone. In contrast, we found that the purely time dependent heuristic policy using a perfect biomarker resulted in an incremental QALY gain of 0.10. This suggests there is the potential to improve the purely time dependent policies using cystoscopy by adding a biomarker.

We compared the heuristic-bmk policy using each of the six biomarkers approved by FDA. We observed that the heuristic-bmk policies resulted in a large variance of the average of biomarker tests, from 7.41 to 16.13; but they had similar average number of cystoscopies, from 7.03 to 7.67. We found that the waiting periods after a negative biomarker result, b_t, are most

affected by the sensitivity of the biomarker compared with the specificity of the biomarker; we also found that the waiting periods after a negative triggered cystoscopic result, c_t, are most affected by the specificity of the biomarker compared with the sensitivity of the biomarker.

5.6 Conclusions

Although there is a significant amount of research on the development of new biomarker tests, there is a lack of studies to show if urine based biomarkers can improve bladder cancer surveillance, and if so how they can be integrated into a surveillance protocol. Our results show that using a biomarker based surveillance protocol has very little potential to improve the optimal policy using cystoscopy alone in terms of QALY gain for low risk bladder cancer patients. It is worth noting that this results is based on a base case disutility of 0.003 for cystoscopy. For larger disutilities, such as may be reasonable for certain patients,the benefit of the biomarker may be larger.

We also developed easy-to-implement heuristic policies to investigate the practical potential for applying a a protocol based on a biomarker to direct the frequency of tests based on the time since the start of surveillance. We found that using the easy-to-implement policies gained from simulating the optimal POMDP solutions can result in significant QALY gain over the EAU and AUA guidelines. We investigated some general rules for using a biomarker based on its sensitivity and specificity. Our results suggest that given a biomarker with higher sensitivity, we should wait a longer period of time before the next biomarker test after a negative biomarker result; given a biomarker with higher specificity, we should wait a shorter period of time before the next biomarker test after a negative triggered cystoscopic result. We also observed that the purely time dependent heuristic policy using a biomarker based surveillance protocol has significant potential to improve the purely time dependent heuristic policy using cystoscopy alone.

Chapter 6

Conclusions

6.1 Summary

This dissertation investigated the design of optimal surveillance policies for low risk bladder cancer patients. First, we used Monte-Carlo simulation to compare the published EAU and AUA guidelines. Next, we extended our simulation model to a POMDP model to investigate the optimal cystoscopy-based surveillance policy. Finally, we extended the POMDP model to include new urine based biomarker tests in the surveillance policy. We analyzed the incremental benefit of biomarker tests for improving the optimal surveillance policies and we investigated easy-to-implement heuristic surveillance policies.

Chapter 2 provided some background on bladder cancer and motivation for studying low risk bladder cancer surveillance. We also provided a methodological review of the POMDP literature, and some recent applications of POMDPs to medical decision making. In Chapter 3 we described the partially observable Markov model based on states that define patient risk levels associated with recurrence and progression of bladder cancer. Monte-Carlo sampling was used to generate sample paths for each surveillance policy to estimate expected QALYs. Published AUA and EAU guidelines were compared with alternative policies based on expected QALYs over a base case patient's lifetime. Finally, sensitivity analysis on model input parameters was presented.

In Chapter 4, we extended the Markov model to a POMDP to study the optimal surveillance policy that maximizes expected QALYs. Optimal policies were computed using the incremental pruning algorithm. We compared the optimal policy to the published AUA and EAU guidelines studied in Chapter 3. We also performed sensitivity analysis on the optimal policy with respect to model input parameters.

In Chapter 5 we extended the POMDP model of Chapter 4 by incorporating a urine based biomarker test. In addition we investigated some easy-to-implement surveillance schedules

obtained with heuristic methods by simulating the optimal policy among patients cohorts with similar heath profile. We analyzed the incremental benefits of biomarkers in terms of the incremental QALY gain, by comparing an easy-to-implement heuristic policy using a biomarker based protocol with an easy-to-implement heuristic policy using cystoscopy alone. We also performed one-way sensitivity analysis and compared the heuristic policies using alternative biomarkers with varying sensitivity and specificity.

6.2 Conclusions

Our study in Chapter 3 suggests that the EAU guideline yields higher expected QALYs but also higher life-long progression probabilities than the AUA policy. We found that patients with lower all other cause mortality should undergo more intensive surveillance. Our sensitivity analysis showed that patients should undergo more intensive surveillance if the disutility of cystoscopy is reduced. Based on the results in Chapter 3, we conclude that patient specific factors such as the presence of comorbidity, or perception of utility loss from cystoscopy, should be considered in determining the best surveillance policy for an individual patient. Although the differences among polices on the basis of QALYs is relatively small, our bi-criteria analysis revealed there are significant differences among policies in the number of cystoscopies. We observed that the number of cystoscopies over a patient's lifetime ranged from 4.13 for strategy D_{12} to 13.76 for the AUA strategy. We found that no one policy dominated another, i.e., all policies were on the efficient frontier. We observed that the EAU policy resulted in nearly half of the number of cystoscopies with a reduction in the relative progression risk of 17% and a reduction in the absolute progression risk of 0.4%. The large variation in the number of cystoscopies among surveillance policies, particularly in the context of the very low background rate of progression to invasive cancer in this population, underscores the importance of understanding the quality of life impact of this management practice on patients.

In Chapter 4, we used a POMDP model to investigate the optimal surveillance policies that maximize a patient's expected QALYs. From the base case scenario we observed that the optimal policies can result in a 0.13 QALY gain and 0.16 life year gain compared with the EAU and AUA guidelines, respectively. We performed sensitivity analysis to determine which model parameters most affect the optimal policy. We observe that the frequency of the optimal surveillance policy is highly sensitive to the disutility of cystoscopy. Changing the disutility of cystoscopy from 0.05 to 0.0015 results in an increase of number of cystoscopies from 0 to 11.20 on average. This indicates that a patient's personal preference for cystoscopy, i.e., the extent to which it affects his or her quality of life, should be considered in the design the optimal surveillance policy. We also found that the optimal surveillance policy can be significantly affected by other cause mortality in terms of incremental QALY gain and incremental life years

gain over the EAU and AUA guidelines, as well as the number of cystoscopies. Changing other cause mortality from 150% of the base case value to 50% of the base case value results in an increase of QALYs (life years) gained from 0.02 (0.02) to 0.35 (0.43) and an increase of number of cystoscopies from 5.20 to 14.00. This indicates that patients with lower other cause mortality may benefit more from the optimal policy and these patients should generally follow more intensive surveillance than the average patient.

In Chapter 5, we extended the POMDP model of Chapter 4 to include urine based biomarker tests for surveillance. To investigate the possible benefit of using a biomarker to direct the frequency of diagnostic tests, we evaluated two different surveillance protocols (with and without using a biomarker). We observed that the QALY gain from the optimal policy using a biomarker based surveillance protocol is not statistically different from the optimal policy using cystoscopy alone for our base case numerical experiments. A further comparison between the optimal policy using a perfect biomarker (100% sensitivity and 100% specificity) and the optimal policy using cystoscopy alone shows that the potential to use a biomarker based protocol to improve on QALYs is very limited.

Based on the sensitivity analysis in Chapter 5, we found that the optimal surveillance frequency using a biomarker based protocol is highly sensitive to the specificity of the biomarker. Patients should undergo more biomarker tests if a biomarker with higher specificity is used for surveillance. Specifically, varying the specificity of the biomarker from 47% (80% of the base case) to 71% (120% of the base case) resulted an increase of the average biomarker tests from 7.67 to 12.37. We also found that the disutility of cystoscopy affected the optimal policy using a biomarker based surveillance protocol. Specifically, varying the disutility of cystoscopy from 0.05 to 0.015 resulted in an increase in the number of biomarker tests from 0 to 14.96 on average. Therefore, the optimal strategy using a biomarker based surveillance protocol should also be influenced by the individual patient's perception of the disutility of cystoscopy. This finding is similar to our finding in Chapter 4 that the optimal policy using cystoscopy alone should be influenced by a patient's personal perception of the disutility of cystoscopy.

Based on the sensitivity analysis on the heuristic policy using cystoscopy alone, heuristic-cyst, we found that the intervals of a heuristic-cyst policy can be affected by the sensitivity of cystoscopy, which is consistent with our previous finding in Chapter 3. Our results suggest patients should wait for a longer time period given a negative result from a cystoscopy with higher sensitivity is observed. Based on the sensitivity analysis on the heuristic policy using a biomarker based surveillance protocol, heuristic-bmk, we observed that the waiting period after a negative biomarker test is positively correlated with both the sensitivity and the specificity of the biomarker; we also observed that the waiting period after a negative triggered cystoscopic result is negatively correlated with both the sensitivity and the specificity of the biomarker.

We further compared the heuristic-bmk policy using each of six biomarkers approved by

FDA. We observed that for all six biomarkers, except NMP22, the heuristic-bmk policy resulted in significant incremental QALY gain over the heuristic-cyst policy. Specifically, the heuristic-bmk policies using each of the other five biomarkers (which all have higher specificity (on low and intermediate risk NMIBC) and higher sensitivity than NMP22) resulted in a 11.25 QALY gain compared with 11.21 QALY gain from the heuristic-cyst policy using a cystoscopy alone. Furthermore, we found that the heuristic policy using a perfect biomarker (perfect sensitivity and perfect specificity) resulted in an incremental QALY gain of 0.08 compared with the heuristic policy using cystoscopy alone. These findings indicate that the purely time dependent heuristic policy using the biomarker based surveillance protocol has the potential to significantly improve the purely time dependent heuristic policy based on using cystoscopy alone. We also observed that the heuristic-bmk policies resulted in large variation in the average of biomarker tests, from 7.41 to 16.13; but similar average number of cystoscopies, from 7.03 to 7.67. We found that the waiting periods after a negative biomarker result can be most affected by the sensitivity of the biomarker compared with the specificity of the biomarker; we also found that the waiting periods after a negative triggered cystoscopic result can be most affected by the specificity of the biomarker compared with the sensitivity of the biomarker.

6.3 Limitations

There were some limitations to our study. First, we made several assumptions in the development of our simulation model and POMDP models regarding the timing, adherence, and duration of treatment. Among these, we assumed that treatment is immediately triggered after detection of disease, and patients' adherence to treatment is perfect. We also assumed that treatments will always be done within a decision epoch (e.g. one month), while in reality BCG and chemotherapy may last half year or even longer. Second, estimates of bladder cancer progression rates were not available from the literature, therefore we calculated the implied monthly progression rates using the simulation model of Chapter 3 by minimizing the difference between the model output of 5 year progression rates with that in the EORTC risk table. Finally, model parameter estimates come from various sources in the literature. For example, we used an estimate of disutility for cystoscopy from a study by Kulkarni et al. [40], which in turn drew from studies of disutility of other moderately invasive procedures, not specifically cystoscopy.

6.4 Future Research Opportunities

Low risk bladder cancer patients have a much lower risk of disease progression and cancer-specific mortality compared than intermediate risk and high risk patients. Nevertheless, current

guidelines do not distinguish patients on the basis of risk. The results of this thesis suggest some further research opportunities to more completely characterize the quality of life impact of frequent, repetitive, invasive procedures in this context, and possibilities for improving surveillance policies using new urine based biomarker tests for low risk bladder cancer patients. First, due to lack of empirical estimates of cystoscopy disutility, future studies should investigate individual patient preferences for cystoscopy to determine how variation in disutility among patients may influence optimal personalized surveillance protocols. Second, since it is desirable to develop easy-to-implement surveillance policies that can improve published guidelines, future work could extend the heuristic policies in Chapter 5 to investigate the possibility of developing better (closer to optimal) policies with easy-to-implement structures. Third, our study in Chapter 5 was limited to two particular protocols. Relaxing the biomarker-based protocol assumption that a cystoscopy is always triggered after a positive biomarker test could lead to improvement in surveillance polices. Finally, given that bladder cancer has been characterized among cancers as having the highest cost per patient from diagnosis to death, future work should include the economic evaluation for using biomarkers in surveillance. The work presented in this thesis provides a foundation for these future studies.

REFERENCES

[1] SEER Cancer Statistics Review 1975-2008 National Cancer Institute SEER Cancer Statistics Review 1975-2008 National Cancer Institute. Technical report, 2011.

[2] Jae-Hyeon Ahn and John C. Hornberger. Involving Patients in the Cadaveric Kidney Transplant Allocation Process: A Decision-Theoretic Perspective. *Management Science*, 42(5):629–641, 1996.

[3] Oguzhan Alagoz, Lisa M. Maillart, Andrew J. Schaefer, and Mark S. Robert. The Optimal Timing of Living-Donor Liver Transplantation. *Management Science*, 50(10):1420–1430, 2004.

[4] M. Babjuk, W. Oosterlinck, R. Sylvester, E. Kaasinen, and J. Palou. Guidelines on TaT1(Non-muscle invasive) Bladder Cancer. Technical report, European Association of Urology, 2009.

[5] M. F. Botteman, C. L. Pashos, R. S. Hauser, B. L. Laskin, and A. Redaelli. Quality of life aspects of bladder cancer: A review of the literature. *Quality of life research : an international journal of quality of life aspects of treatment, care and rehabilitation*, 12:675–688, September 2003.

[6] Marc F. Botteman, Chris L. Pashos, Alberto Redaelli, Benjamin Laskin, and Robert Hauser. The Health Economics of Bladder Cancer: A Comprehensive Review of the Published Literature. *Pharmacoeconomics*, 21(18):1315–30, January 2003.

[7] Peter Boyle and Bernard Levin. World Cancer Report 2008. Technical report, International Agency for Research on Cancer, Lyon, France, 2008.

[8] Anthony Cassandra, Michael L. Littman, and Nevin L. Zhang. Incremental Pruning : A Simple, Fast, Exact Method for Partially Observable Markov Decision Processes. In *Proceedings Thirteenth Annual Conference on Uncertainty in Artificial Intelligence*, number Cassandra1996, San Francisco, CA, 1997.

[9] Anthony R. Cassandra. A Survey of POMDP Applications. In *AAAI Fall Symposium: Planning with POMDPs*, Orlando, FL, 1998.

[10] Anthony R. Cassandra. *Exact and Approximate Algorithms for Partially Observable Markov Decision Processes.* Phd dissertation, Brown Univeristy, 1998.

[11] R. Chahal, S. K. Sundaram, R. Iddenden, D. F. Forman, P. M. T. Weston, and S. C. W. Harrison. A Study of the Morbidity, Mortality and Long-Term Survival Following Radical Cystectomy and Radical Radiotherapy in the Treatment of Invasive Bladder Cancer in Yorkshire. *European Urology*, 43:246–257, 2003.

[12] H.T. Cheng. *Algorithms for partially observable Markov decision processes.* Phd dissertation, University of British Columbia, Vancouver, BC, 1988.

[13] Liang Cheng, Shaobo Zhang, Gregory T Maclennan, Sean R Williamson, Antonio Lopez-Beltran, and Rodolfo Montironi. Bladder cancer: translating molecular genetic insights into clinical practice. *Human Pathology*, 42(4):455–481, 2011.

[14] Marc Colombel, Mark Soloway, Hideyuki Akaza, Andreas Bo, Joan Palou, Roger Buckley, Donald Lamm, Maurizio Brausi, J. Alfred Witjes, and Raj Persad. Epidemiology, Staging, Grading, and Risk Stratification of Bladder Cancer. *European Urology*, 7:618–626, 2008.

[15] R. R. de Vries, J. A. Nieuwenhuijzen, A. Vincent, H. van Tinteren, and S. Horenblas. Survival after Cystectomy for Invasive Bladder Cancer. *European Journal of Surgical Oncology*, 36:292–297, 2010.

[16] S Donat. Evaluation and follow-up strategies for superficial bladder cancer. *Urologic Clinics of North America*, 30(4):765–776, 2003.

[17] James E. Eckles. *Optimum replacement of stochastically failing systems*. Phd dissertation, Stanford University, Stanford, CA, 1966.

[18] Nader M. Fahmy, Salaheddin Mahmud, and Armen G. Aprikian. Delay in the Surgical Treatment of Bladder Cancer and Survival: Systematic Review of the Literature. *European Urology*, 50:1176–1182, 2006.

[19] Ubirajara Ferreira, Wagner Eduardo Matheus, Renato Nardi Pedro, Carlos Alturo Levi D'Ancona, Leonardo Oliveira Reis, Rafael Mamprin Stopiglia, Fernandes Denardi, Nelson Rodrigues Jr. Netto, Stenio De Cassio Zequi, Francisco Paulo Da Fonseca, Ademar Lopes, Gustavo Cardoso Guimaraes, Roni De Carvalho Fernandes, and Marjo Deninson Cardenuto Perez. Primary Invasive versus Progressive Invasive Transitional Cell Bladder Cancer: Multicentric Study of Overall Survival Rate. *Urologia Internationalis*, 79:200–203, 2007.

[20] A. Herbert Fritsche, H. Barton Grossman, Seth P. Lerner, and Ihor Sawczuk. National Academy of Clinical Biochemistry Guidelines for the Use of Tumor Markers in Bladder Cancer. *Cancer*, 2008.

[21] Ofer N. Gofrit, Dov Pode, Adi Lazar, Ran Katz, and Amos Shapiro. Watchful Waiting Policy in Recurrent Ta G1 Bladder Tumors. *European Urology*, 49:303–307, 2006.

[22] H. Barton Grossman, Leonard Gomella, Yves Fradet, Alvaro Morales, Joseph Presti, Chad Ritenour, Unyime Nseyo, and Michael J Droller. A Phase III, Multicenter Comparison of hexaminolevulinate Fluorescence Cystoscopy and White Light Cystoscopy for the Detection of Superficial Papillary Lesions in Patients With Bladder Cancer. *The Journal of Urology*, 178(7):62–67, 2007.

[23] H. Barton Grossman, Mark Soloway, Edward Messing, Giora Katz, Barry Stein, Vahan Kassabian, and Yu Shen. Surveillance for Recurrent Bladder Cancer Using a Point-of-Care Proteomic Assay. *JAMA*, 295(3):299–305, 2006.

[24] M. Craig Hall, Sam S. Chang, Guido Dalbagni, Raj Som Pruthi, John Derek Seigne, Eila Curlee Skinner, J. Stuart Wolf, and Paul F. Schellhammer. Guideline for the Management of Nonmuscle Invasive Bladder Cancer (Stages Ta, T1, and Tis): 2007 update. *The Journal of Urology*, 178(6):2314–30, December 2007.

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

[26] Melonie Heron, Donna L Hoyert, Sherry L Murphy, Jiaquan Xu, Kenneth D. Kochanek, and Betzaida Tejada-Vera. Deaths: Final Data for 2006. *National Vital Statistics Reports*, 57(14), 2009.

[27] Harry W Herr. Natural history of superficial bladder tumors: 10- to 20-year follow-up of treated patients. *World Journal of Urology*, 15(2):84–88, January 1997.

[28] Harry W Herr, S Machele Donat, and Victor E Reuter. Management of Low Grade Papillary Bladder Tumors. *The Journal of Urology*, 178(October):1201–1205, 2007.

[29] Harry W Herr and Marisa Schneider. Outpatient flexible cystoscopy in men: a randomized study of patient tolerance. *The Journal of Urology*, 165(June):1971–1972, 2001.

[30] Sten Holmang, Hans Hedelin, Claes Anderstrom, Erik Holmberg, Christer Busch, and Sonny L. Johansson. RECURRENCE AND PROGRESSION IN LOW GRADE PAPILLARY UROTHELIAL TUMORS. *The Journal of Urology*, 162(3):702–707, 1999.

[31] MJ Horner, LAG Ries, and M Krapcho. SEER Cancer Statistics Review 1975-2006. 2006.

[32] Ronald A. Howard. *Dynamic Probabilistic Systems*. Weiley, New York, 1973.

[33] Maryska L.G. Janssen-heijnen, Adam Gondos, Freddie Bray, Timo Hakulinen, David H. Brewster, Hermann Brenner, and Jan-willem W. Coebergh. Clinical Relevance of Conditional Survival of Cancer Patients in Europe: Age-Specific Analyses of 13 Cancers. *Journal of Clinical Oncology*, (28), 2010.

[34] Ahmedin Jemal, Rebecca Siegel, Jiaquan Xu, and Elizabeth Ward. Cancer Statistics, 2010. *CA: A Cancer Journal for Clinicians*, 60:277–300, 2010.

[35] Pack Leslie Kaebling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101:99–134, 1998.

[36] D. L. Kent, R. Shachter, H. C. Sox, L. D. Shortliffe, S. Moynihan, and F. M. Torti. Efficient Scheduling of Cystoscopies in Monitoring for Recurrent Bladder Cancer. *Medical Decision Making*, 9(1):26–37, February 1989.

[37] Daniel L. Kent, Robert A. Nease, Harold C. Sox, Linda D. Shortliffe, and Ross Shachter. Eualuation of Nonlinear Optimization for Scheduling of follow-up cystoscopies to Detect Recurrent Bladder Cancer. *Medical Decision Making*, 11:240–248, 1991.

[38] Ziya Kirkali, Theresa Chan, Murugesan Manoharan, Ferran Algaba, Christer Busch, Liang Cheng, Lambertus Kiemeney, Martin Kriegmair, R. Montironi, William M. Murphy, Isabell A Sesterhenn, Masaaki Tachibana, and Jeff Weider. Bladder Cancer: Epidemiology, Staging and Grading, and Diagnosis. *Urology*, 66(Suppl 6A):4–34, 2005.

[39] Fumitaka Koga, Soichiro Yoshida, Satoru Kawakami, Yukio Kageyama, Minato Yokoyama, Kazutaka Saito, Yasuhisa Fujii, Tsuyoshi Kobayashi, and Kazunori Kihara. Low-Dose Chemoradiotherapy Followed by Partial or Radical Cystectomy Against Muscle-Invasive Bladder Cancer: An Inten-to-Treat Survival Analysis. *Oncology*, 72:384–388, 2008.

[40] Girish S. Kulkarni, Shabbir M. H. Alibhai, Antonio Finelli, Neil E. Fleshner, Michael A. S. Jewett, Steven R Lopushinsky, and Ahmed M. Bayoumi. Cost-Effectiveness Analysis of Immediate Radical Cystectomy Versus Intravesical Bacillus Calmette-Guerin Therapy for High-Risk, High-Grade (T1G3) Bladder Cancer. *Cancer*, 115(23):5450–5459, December 2009.

[41] Girish S. Kulkarni, Antonio Finelli, Neil E. Fleshner, Michael A. S. Jewett, Steven R. Lopushinsky, and Shabbir M. H. Alibhai. Optimal Management of High-Risk T1G3 Bladder Cancer: A Decision Analysis. *PLoS Medicine*, 4(9):1538–1549, September 2007.

[42] Moshe Leshno, Zamir Halpern, and Nadir Arber. Cost-Effectiveness of Colorectal Cancer Screening in the Average Risk Population. *Health Care Management Science*, 6:165–174, August 2003.

[43] Michael L. Littman. The Witness Algorithm: Solving Partially Observable Markov Decision Processes. Technical Report December, Department of Computer Science, Providence, Rhode Island, 1994.

[44] Michael L. Littman, Anthony R. Cassandra, and Leslie Pack Kaelbling. Effcient dynamic-programming updates in partially observable Markov decision processes. Technical Report November, Brown University, Providence, Rhode Island, 1995.

[45] William S. Lovejoy. A Survey of Algorithmic Methods for Partially Observed Markov Decision Processes. *Annals of Operations Research*, 28:47–66, 1991.

[46] William S. Lovejoy. Computationally Feasible Bounds for Partially Observed Markov Decision Processes. *Operations Research*, 39(1):162–175, 1991.

[47] Lars Lund, Jacob Jacobsen, Peter Clark, Michael Borre, and Mette Nø rgaard. Impact of Comorbidity on Survival of Invasive Bladder Cancer Patients, 1996-2007: A Danish Population-based Cohort Study. *Journal of Urology*, 75:393–398, 2010.

[48] Stephan Madersbacher, Werner Hochreiter, Fiona Burkhard, George N. Thalmann, Regula Markwalder, and Urs E. Studer. Radical Cystectomy for Bladder Cancer Today-A Homogeneous Series Without Neoadjuvant Therapy. *Journal of Clinical Oncology*, 21(4):690–696, 2003.

[49] Lisa M. Maillart, Julie Simmons Ivy, Scott Ransom, and Kathleen Diehl. Assessing Dynamic Breast Cancer Screening Policies. *Operations Research*, 56(6):1411–1427, 2008.

[50] Maria Elena Martinez, John A. Baron, David A. Liberman, Arthur Schatzkin, Elaine Lanza, Sidney J. Winawer, ANN G. Zauber, Ruiyun Jiang, Dennis J Ahnen, John H Bond, Timothy R Church, Douglas J Robertson, Stephanie A Smith-Warner, Elizabeth T Jacobs, David S Alberts, and E Robert Greenberg. A Pooled Analysis of Advanced Colorectal

Neoplasia Diagnoses After Colonoscopic Polypectomy. *Gastroenterology*, 136(3):832–841, 2009.

[51] George E. Monahan. A Survey of Partially Observable Markov Decision Processes: Theory, Model, and Algorithms. *Management Science*, 28(1):1–16, 1982.

[52] James E. Montie, Peter E. Clark, Mario A. Eisenberger, Rizk EI-Galley, Richard E. Greenberg, Harry W. Herr, Gary R. Hudes, Deborah A. Kuban, Timothy M. Kuzel, Paul H. Lange, Subodh M. Lele, Jeffrey Michalski, Anthony Patterson, Kamal S. Pohar, Jerome P. Richie, Wade J. Sexton, William U. Shipley, Eric J. Small, Donald L. Trump, Philip J. Walther, and Timothy G. Wilson. Practice Guidelines in Oncology - Bladder Cancer Including Upper Tract Tumors and Urothelial Carcinoma of the Prostate. *Practice Guidelines in Oncology*, 2010.

[53] G. Mowatt, S. Zhu, M. Kilonzo, C. Boachie, C. Fraser, T. R. L. Griffiths, J. N. Dow, G. Nabi, J. Cook, and L. Vale. Sytematic reivew of the clinical effectiveness and cost-effectiveness of photodynamic diagnosis and urine biomarkers (FISH, ImmunoCyt, NMP22) and cytology for the detection and follow-up of bladder cancer. *Health Technology Assessment*, 14(4), 2010.

[54] Willem Oosterlinck, Eduardo Solsona, Hideyuki Akaza, Christer Busch, Peter J. Goebell, Per-Uno Malmström, Haluk Ozen, and Paul Sved. Low-Grade Ta (Noninvasive) Urothelial Carcinoma of The Bladder. *Urology*, 66(Suppl 6A):75–89, 2005.

[55] Raj Persad, Donald Lamm, Maurizio Brausi, Mark Soloway, Joan Palou, Andreas Bo, Marc Colombel, Hideyuki Akaza, Roger Buckley, and Alfred Witjes. Current Approaches to the Management of Non-Muscle Invasive Bladder Cancer: Comparison of Current Guidelines and Recommendations. *European Urology Supplements*, 7:637–650, 2008.

[56] Pascal Poupart. *Exploiting Structure to Efficiently Solve Large Scale Partially Observable Markov Decision Process*. Phd dissertation, University of Toronto, 2005.

[57] Raj S. Pruthi, Nathan Baldwin, Vishal Bhalani, and Eric M. Wallen. Conservative Management of Low Risk Superficial Bladder Tumors. *The Journal of Urology*, 179(1):87–90, 2008.

[58] M. Sant, T. Aareleid, F. Berrino, M. Bielska Lasota, P. M. Carli, J. Faivre, P. Grosclaude, G. Hedelin, T. Matsuda, H. Moller, T. Moller, A. Verdecchia, R. Capocaccia, G. Gatta, A. Micheli, M. Santaquilani, P. Roazzi, and D. Lisi. EUROCARE-3: survival of cancer patients diagnosed 1990-94–results and commentary. *Annals of Oncology*, 14(Supplement 5):v61–v118, 2003.

[59] Andrew J. Schaefer, Matthew D. Bailey, Steven M. Shechter, and Mark S. Roberts. Modeling Dedical Treatment Using Markov Decision Processes. In *Operations Research and Health Care: A Handbook of Methods and Applications*, chapter 23, pages 598–616. Springer, New York, 2005.

[60] Deborah Schrag, Lillian J. Hsieh, Farhang Rabbani, Peter B. Bach, Harry Herr, and Colin B Begg. Adherence to Surveillance Among Patients With Superficial Bladder Cancer. *Journal of the National Cancer Institute*, 95(8):588–597, 2003.

[61] Shahrokh F. Shariat, Pierre I. Karakiewicz, Ganesh S. Palapattu, Yair Lotan, Craig G. Rogers, Gilad E. Amiel, Amnon Vazina, Amit Gupta, Patrick J. Bastian, Arthur I. Sagalowsky, Mark P. Schoenberg, and Seth P. Lerner. Outcomes of Radical Cystectomy for Transitional Cell Carcinoma of the Bladder: A Contemporary Series From the Bladder Cancer Research Consortium. *The Journal of Urology*, 176(6):2414–2422, 2006.

[62] Richard D. Smallwood and Edward J. Sondik. The Optimal Control of Partially Observable Markov Processes Over a Fnite Horizon. *Operations Research*, 21(5):1071–1088, 1973.

[63] Richard D. Smallwood, Edward J. Sondik, and Fred L. Offensend. Toward an Integrated Methodology for Analysis of Health-Care Systems. *Operations Research*, 19(6):1300–1322, 1971.

[64] Mark S. Soloway, Darren S. Bruck, and Sandy S. Kim. Expectant Manangement Of Small, Recurrent, Noninvasive Papillary Bladder Tumors. *The Journal of Urology*, 170(8):438 – 441, 2003.

[65] John P. Stein, Gary Lieskovsky, Richard Cote, Susan Groshen, An-chen Feng, Stuart Boyd, Eila Skinner, Bernard Bochner, Duriayai Thangathurai, Maged Mikhail, Derek Raghavan, and Donald G. Skinner. Radical Cystectomy in the Treatment of Invasive Bladder Cancer: Long-Term Results in 1,054 Patients. *Journal of Clinical Oncology*, 19(3):666–675, 2001.

[66] Richard J. Sylvester, Adrian P. M. van der Meijden, Willem Oosterlinck, J. Alfred Witjes, Christian Bouffioux, Louis Denis, Donald W. W. Newling, and Karlheinz Kurth. Predicting Recurrence and Progression in Individual Patients with Stage Ta T1 Bladder Cancer Using EORTC Risk Tables: A Combined Analysis of 2596 Patients from Seven EORTC Trials. *European Urology*, 49:466–477, March 2006.

[67] Bas W. G. van Rhijn, Henk G. van Der Poel, and Theo H. van Der Kwast. European Urology Urine Markers for Bladder Cancer Surveillance: A Systematic Review. *European Urology*, 47:736–748, 2005.

[68] Olaf P. J. Vrooman and J. Alfred Witjes. Urinary Markers in Bladder Cancer. *Urology*, 53:909–916, 2008.

[69] Chelsea C. White. A Survey of Solution Techniques for The Partially Observed Markov Decision Process. *Annals of Operations Research*, 32:215–230, 1991.

[70] Erika M Wolff, Gangning Liang, and Peter A Jones. Mechanisms of Disease: genetic and epigenetic alterations that drive bladder cancer. *Nature Clinical Practice Urology*, 2(10):502–510, 2005.

[71] Jingyu Zhang, Brian Denton, H. Balasubramanian, B. Inman, and N. Shah. Estimating the True Value of PSA Tests for Prostate Cancer Detection. *Working paper*, 2009.

[72] Jingyu Zhang, Brian Denton, and Brant Inman. Optimization of PSA-Based Screening Decisions for Prostate Cancer Detection. *Working paper*, 2009.

[73] Nevin L. Zhang and Wenju Liu. Planning in Stochastic Domains: Problem Characteristics and Approximation. Technical report, The Hong Kong University of Science and Technology, Hong Kong, 1996.