

**Nursing Workforce Planning and Radiation Therapy  
Treatment Decision Making: Two Healthcare Operations  
Research Applications**

by

**Marisel Sofia Lavieri**

B.A., The University of Florida, 2002  
B.Sc., The University of Florida, 2002  
M.Sc., The University of British Columbia, 2004

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

**DOCTOR OF PHILOSOPHY**

in

**The Faculty of Graduate Studies**

**(Business Administration)**

**THE UNIVERSITY OF BRITISH COLUMBIA  
(Vancouver)**

**September 2009**

**© Marisel Sofia Lavieri, 2009**

## ABSTRACT

This thesis discusses two applications of operations research to healthcare: nursing workforce planning and radiation therapy treatment decision making. The first application describes a linear programming based hierarchical planning tool that determines the optimal number of nurses to train, promote to managerial levels and recruit over a 20 year planning horizon to achieve nursing and managerial targets. The model is based on the age dynamics and attrition rates of the nursing workforce. The tool has been developed to assist policy makers in planning the British Columbia registered nurses workforce. Its simplicity of use makes it ideal for scenario and “What-If” analyses.

The second application presents a novel approach to model individual disease progression of prostate cancer patients who receive neoadjuvant hormone therapy before radiation therapy. The model is used to help clinicians determine when to initiate radiation therapy based on a patient’s prostate specific antigen (PSA) dynamics. Each patient’s PSA dynamics is modeled by a log quadratic curve. Prior distributions for the curve parameters are obtained from population characteristics. The distribution of the time of the PSA nadir (which might be linked to maximal tumor regression) is derived from an approximation of the ratio of two correlated normal random variables. Using a dynamic Kalman filter model, the parameter estimates are updated as new patient specific information becomes available. Clustering is incorporated to improve our prior estimates of the curve parameters. The model trades off the risk of beginning radiation therapy too soon, before hormone therapy has achieved its maximum effect, against waiting too long to start therapy after there has been a potential increase in the number of tumor cells resistant to the treatment. We illustrate and validate our modeling approach by comparing clinically implementable policies (cumulative probability policy and threshold probability policy) on a cohort of prostate cancer patients, and show that our approach outperforms the current protocol by identifying earlier when radiation therapy should start for each patient.

While both applications involve very different approaches, they incorporate dynamic decision making in the field of healthcare. A deeper knowledge of the potential of the field is achieved by understanding the challenges faced and methodology used to guide decisions on a policy level as well as on a patient-specific level.

# TABLE OF CONTENTS

Abstract.....	ii
Table of Contents .....	iii
List of Tables .....	v
List of Figures .....	vi
Acknowledgements.....	viii
Dedication .....	ix
Co-Authorship Statement .....	x
Chapter 1: Introduction.....	1
Chapter 2: Optimal Planning of the British Columbia Registered Nurses' Workforce.....	4
2.1 Nursing in British Columbia.....	6
2.2 The Model .....	8
2.2.1 Model Assumptions .....	9
2.2.2 Decision Variables.....	11
2.2.3 Data Formulation .....	14
2.2.4 Model Formulation .....	17
2.3 Implementation.....	20
2.3.1 Scenario Analysis.....	21
2.3.2 Results.....	24
2.4 Discussion .....	28
2.5 Conclusions and Remarks .....	31
Chapter 3: A Patient Specific Model for Deciding When to Start Radiation Therapy for Prostate Cancer Patients Based on Their PSA Dynamics.....	33
3.1 Methods .....	36
3.1.1 Modeling PSA Dynamics.....	36
3.1.2 Estimating the Prior Distribution of the Curve Parameters.....	38
3.1.3 Modeling the Distribution of the Nadir .....	39

3.1.4 Updating the Parameter Estimation Using a State Space Model .....	43
3.1.5 Clustering.....	44
3.2 Implementation and Data Analysis .....	47
3.2.1 Data Analysis.....	48
3.2.2 Clinical Decision Making .....	54
3.3 Conclusions and Remarks .....	61
Chapter 4: Validation of a PSA Dynamics Model and Nadir Prediction Methods.....	63
4.1 Methodology .....	63
4.2 Data Description.....	66
4.3 Model Validation .....	68
4.4 Policy Implications.....	76
4.5 Conclusion and Remarks .....	79
Chapter 5: Conclusions, Future Challenges and Final Remarks .....	81
5.1 Main Conclusions .....	81
5.2 Future Challenges.....	83
5.2.1 The Workforce Planning Model.....	83
5.2.2 The Prostate Cancer Model .....	85
5.3 Concluding Remarks.....	86
References.....	88

## LIST OF TABLES

Table 2.1. BC's employed registered nurses' demographics and population statistics, 1997-2005....	7
Table 2.2. Notation of inputs .....	15
Table 3.1. Model notation. ....	36
Table 3.2. Average cluster probabilities given most likely cluster membership.....	50
Table 3.3. Model comparisons based on the mean absolute difference and maximum absolute difference between each policy's RT start time and the estimated <i>tnadir<sub>i</sub></i> for all patients. ....	59
Table 4.1. Number of patients in the validation database. ....	67
Table 4.2. Length of hormone treatment for patients in the validation database (Average Number of Days [Minimum, Maximum]).....	68
Table 4.3. Time (in days) between PSA readings for patients in the validation database (Average, [Minimum, Maximum]). ....	68
Table 4.4. Proportion of not censored patients in the validation database.....	70
Table 4.5. Proportion of patients for whom the estimated nadirs and the current protocol fall within the interval assumed to contain the PSA nadir .....	72
Table 4.6. Comparison of the estimated PSA nadir after the patient has had one, two or three readings to the current protocol based on the proportion that falls within the interval assumed to contain the patient's PSA nadir.....	73
Table 4.7. Estimated time of nadir and minimum PSA value observed by cluster. ....	74

## LIST OF FIGURES

Figure 2.1. Model timeline.....	9
Figure 2.2. Workforce flowchart.....	10
Figure 2.3. Model screenshot.....	21
Figure 2.4. Age distribution of direct care nurses in the first and last years of the planning horizon based on the optimal solution to the model.....	24
Figure 2.5. Number of direct care registered nurses per year .....	25
Figure 2.6. Recruitment policies for various scenarios.....	26
Figure 2.7. Admission policies for various scenarios.....	27
Figure 2.8. Promotion policies for various scenarios .....	28
Figure 2.9. Annual student populations by year of schooling obtained from the optimal solution to the baseline model.....	29
Figure 3.1. Model overview.....	35
Figure 3.2. PSA versus time curve fit for a randomly chosen patient. ....	38
Figure 3.3. Impact of incorporating correlation in the distribution of the time of nadir for a randomly chosen patient. ....	40
Figure 3.4. Main steps in clustering the prior distribution.....	45
Figure 3.5. Distribution of the estimated time to reach the nadir and the minimum PSA value observed. ....	45
Figure 3.6. Steps involved in the curve parameter updating.....	47
Figure 3.7. $\beta_i$ and $\gamma_i$ from the regression curve $Y_{i,t} = \alpha_i + \beta_i t + \gamma_i t^2 + v_i$ calculated for each patient in the dataset.....	49
Figure 3.8. Summary of best BIC as a function of the number $M$ of clusters in the population.....	50
Figure 3.9. Cluster classification of patients based on the estimated time of nadir and minimum PSA value observed. ....	51
Figure 3.10. Regression curves based on prior means obtained by weighting parameters within each cluster.....	52
Figure 3.11. Distribution of the estimated time of nadir among patients after each patient has had 1, 2, 3 or 4 PSA readings (Reading 1, Reading 2, Reading 3 and Reading 4, respectively). ....	53

Figure 3.12. Change of cluster classification based on the number of times the probability of being in each cluster is updated (1, 2, 3 or 4 times, represented as Reading 1, Reading 2, Reading 3, and Reading 4).....	54
Figure 3.13. Comparison of the mean absolute difference between the RT start time based on each policy and the estimated time of nadir for each patient in the population.....	57
Figure 3.14. Comparison of the population distribution of the difference and absolute difference of the estimated time of nadir for each patient to the RT start time of that patient based on the 85% cumulative probability policy B, 15% threshold probability policy D and the current protocol.....	57
Figure 3.15. Sample output obtained from the tool.....	60
Figure 4.1. Illustration of interval of PSA nadir based on whether patient is censored.....	65
Figure 4.2. Proportion of patients in the validation database.....	67
Figure 4.3. Distribution of $R^2$ values for patients that had at least four PSA readings.....	69
Figure 4.4. Comparison of the cumulative distribution of the time of the minimum PSA value observed versus the nadir estimated by fitting a log quadratic curve based on all readings.....	70
Figure 4.5. Comparison of the estimated PSA nadirs to the current protocol based on whether they fall within the interval assumed to contain the true nadir.....	71
Figure 4.6. Distribution of the difference and absolute difference between the estimated time of nadir based on the Kalman Filter model and the nadir estimated from a model fit to all data.....	73
Figure 4.7. Change of cluster classification over time for all patients and patients classified based on whether they had LHRH monotherapy or total androgen blockade until RT .....	75
Figure 4.8. Comparison of the cumulative distribution of the time of treatment under each protocol to the distribution of the nadir estimated based on all readings.....	76
Figure 4.9. Distribution of the difference from the estimated time of nadir and the time of the minimum PSA value observed for the proposed policies and the current protocol.....	77
Figure 4.10. Distribution of the difference from the estimated time of nadir for censored and not censored patients using the proposed policies and the current protocol.....	78

## ACKNOWLEDGEMENTS

I would like to express my gratitude to all of you who have made the completion of this thesis possible:

- Professor Martin L. Puterman: thanks for believing in my potential and for being such a fantastic advisor! It has been a pleasure working under your supervision. I hope to be as good an advisor to my own students in the future.
- Professor Steven Shechter, for taking the time to be part of my thesis committee as well as for all his great advice throughout this program.
- Professor Pamela Ratner, for helping me understand the dynamics of the registered nurses population of British Columbia and for all her input in this thesis. I am very grateful to have you as part of my thesis committee.
- Professor Scott Tyldesley, for all his patience in helping me understand the biology behind the prostate cancer models as well as for all his invaluable comments. Having you as part of my committee truly made a difference in the quality of the patient specific models created.
- CIHR: New Emerging Team Grant - Access to Quality Cancer Care, Natural Sciences and Engineering Research Council of Canada, the Itoko Muraoka Fellowship, the Bonder Scholarship for Applied Operations Research in Health Services and the Mathematics of Information Technology and Complex Systems for their funding support.
- British Columbia Nurses' Union (BCNU), the College of Registered Nurses of British Columbia (CRNBC), BC Ministry of Advanced Education (MAE), the University of British Columbia, School of Nursing and Dr. William J. Morris and the BC Cancer Agency for the provision of data. Any errors in the data analysis or interpretations are the author's.
- Members of the CIHR Team in Operations Research for Improved Cancer Care, for their assistance and helpful comments.
- Members of the Centre for Healthcare Management, for their support.
- Elaine Cho, for all the help provided during the program.
- Sandra Regan (MSN, RN), for being not only a coauthor of the nursing workforce planning chapter, but also a true friend.
- Mostafa Asadi-M, for his help in the creation of the user friendly PSA dynamics tool.
- Luis Felipe: I couldn't have finished this thesis on time without your help!



## DEDICATION

To my parents Luis and Sophie, my brother Luis Felipe, my boyfriend Michael, my grandparents Pepo, Lela, Oma and Opa and all my family and friends in Canada, the United States, Venezuela, Germany and Brazil. I am blessed to have you in my life. This thesis is for you.

## CO-AUTHORSHIP STATEMENT

Versions of Chapter 2 have been published in:

Lavieri, M.S., S. Regan, M.L. Puterman, and P.A. Ratner. 2008. "Using Operations Research to Plan the British Columbia Registered Nurses' Workforce." *Health Care Policy*. 4(2): 113-131. The original publication is available at [www.longwoods.com](http://www.longwoods.com).

Lavieri, M.S. and M.L. Puterman. 2009. "Optimizing Nursing Human Resource Planning in British Columbia." *Health Care Management Science*. 12: 119-128. The original publication is available at [www.springerlink.com](http://www.springerlink.com). DOI: 10.1007/s10729-008-9097-0.

Versions of Chapters 3 and 4 shall be submitted for publication in collaboration with Professor M. L. Puterman and Professor S. Tyldesley

## CHAPTER 1: INTRODUCTION

Providing quality healthcare that is accessible to all members of the population is a challenge faced by governments worldwide. Timely, efficient healthcare delivery not only significantly impacts patients, but also has significant economic implications: healthcare spending represents over 16% of the US gross domestic product (GDP) (Keehan et al. 2008) and over 10% of the Canadian GDP (CIHI 2008). According to the Canadian Institute for Health Information (2008), it is necessary to have “the right number of health care providers, with the right set of skills, in the right settings” in order to meet the growing health care needs of the population. However, how to achieve and maintain appropriate workforce levels and how to improve treatment decisions based on patient specific outcomes are questions left open. This thesis will address these questions in two specific settings; nurse workforce planning and prostate cancer treatment decision making.

Operations research (OR) has been widely used to model healthcare challenges (Brandeau et al. 2004). Models in the healthcare operations research field may be classified into the following broad categories: policy and public health, tactical operations management and disease management.

Policy and public health models concern strategic decisions that are made over a long term horizon. These decisions usually involve a wide variety of stakeholders including governments, regulatory bodies, professional associations, representatives from the private sector and senior health system executives. Problems faced at this level include trading off capacity expansion and long term capital investments (Lovejoy and Li 2002), deciding where to locate hospitals and health units (Daskin and Dean 2004; Verter and Lapierre 2002), modeling population health and risk management (Chick et al. 2002; Hall et al. 1992; Peterson 1996), relating outcomes to economic analysis (Fuloria and Zenios 2001; Kognakorn and Sainfort 2004; So and Tang 2000), and analyzing how factors such as quality of service are impacted by capacity changes (Güne et al. 2004).

Operation management decisions, on the other hand, arise in the provision of health care at hospitals and other clinical units. The decisions usually involve managers and health professionals. They incorporate problems faced on a daily, weekly, or monthly basis such as scheduling specialties and managing waitlists (Santibañez et al. 2005), assigning resources to shifts (Felici and Gentile 2004) and managing inventory (Henderson and Mason 2004).

Finally, disease management models are related to those decisions that clinicians face in the delivery of health care. Decisions might be related to disease monitoring and screening (Baker 1998; Butler 1979; Harper and Jones 2005; Kirch and Klein 1974; Lincoln and Weiss 1964; Özekici and Pliska 1991; Pardalos et al. 2004; Sherlaw-Johnson and Gallivan 2000; Shwartz 1978) diagnosis (Laporte et al. 1998; Mangasarian 1995; Sanchez de Rivera 1980) and treatment planning (Alagoz et al 2000; Ferris, et al 2003; Hu et al 1996; Mehrez and Gafni 1987; Romeijn et al. 2005; Shechter et al 2008).

In this thesis, I investigate healthcare operations research from both a strategic perspective and a clinician's perspective. Both problems involve dynamic decision making made recursively over a finite planning horizon, yet they use very different approaches. In the former case, I formulate a comprehensive, policy based, methodology for health human resource planning that not only focuses on forecasting future staffing requirements but also on how to best meet those requirements through training, recruitment, and promotions. The model is formulated over a 20 year planning horizon (with decisions made once a year), and it is based on the age dynamics and attrition rates of the providers. This is a novel framework that incorporates variables such as leaves of absence, full time equivalence, promotion rules and attrition from educational programs and the profession. It is applied to assist policy makers in planning the British Columbia registered nurses workforce.

At the disease management level, I develop novel patient specific models that help clinicians determine when to start radiotherapy treatment, based on a patient's prostate specific antigen (PSA) dynamics, for prostate cancer patients receiving hormone therapy. Each time a PSA reading is taken, the clinician must decide whether to start radiation therapy or to wait for a next reading. The model trades off the risk of beginning radiation therapy too soon, before hormone therapy has achieved its maximum effect, against waiting too long to start therapy after there has been a potential increase in the number of tumor cells resistant to the treatment. Prior distributions for the patient's PSA dynamics are obtained from population characteristics, and parameter estimates are updated as new patient specific information becomes sequentially available. Two decision rules are evaluated and shown to identify earlier, and with less variability than under current practice, when the best time to start radiation therapy is likely to be reached. The modeling approach is validated on a cohort of prostate cancer patients from the British Columbia Cancer Agency database.

The thesis is organized as follows. Chapter 2 describes the health human resource planning model, its innovative features as well as the methodological challenges encountered. Chapter 3 describes the prostate cancer decision making model and its implications. This model is validated in Chapter 4. Finally, Chapter 5 reviews key findings and points to new research opportunities. Chapters 2 and 3 contain detailed introductions to these specific areas of research.

## CHAPTER 2: OPTIMAL PLANNING OF THE BRITISH COLUMBIA REGISTERED NURSES' WORKFORCE

A 2004 - Canadian Health Services Research Foundation report stated that workforce planning, training and regulation was the dominant immediate and long term issue for health services policy makers, managers, and clinical organizations. Health professionals' staffing levels significantly impact wait times, patient throughput, and the quality and effectiveness of the services provided by the healthcare system (Tomblin Murphy and O'Brien-Pallas 2004).

Decisions related to workforce supply have long term implications. In order to have a sufficient number of professionals to replace those that will retire in the future, new professionals must be educated. However, admitting too many students is as undesirable as a shortage. Not only does it take time to educate each health professional, but training represents a significant societal financial investment. Furthermore, each newly graduated healthcare professional should be able to take up a position in the healthcare system. Balancing the number of educational seats with the immigration of health professionals from other provinces and other countries while accounting for the impact of other factors on the supply of professionals needs to be considered.

While the importance of decisions pertinent to health human resource planning is acknowledged, no unified model has been developed to determine the best long term plan to meet the forecasted need for health professionals in British Columbia (BC). Models that have been developed in the healthcare field tend to focus on either projecting the impact of current policies based on present workforce characteristics and expected entries and attrition (Kazanjian, Brothers and Wong 1986; Kazanjian, Pulcins and Kerluke 1992; Newton and Burske 1998; Abrahams 2005; Archambault 1999); estimating how many health professionals will be needed based on general population demographics and health status (Basu and Liu 2005; Tomblin Murphy 2005); or estimating requirements based on anticipated future costs or financial resources (Dickson 2002). Most approaches rely heavily on supply-side projections, assume current patterns of utilization and seldom include the impact of policy decisions on supply (O'Brien-Pallas et al. 2001).

These models provide useful insights into the supply and demand for health professionals. However, none is sufficient on its own to address the current and future challenges in planning the health workforce.

Workforce planning has been widely modeled in the operations research literature. The models address two levels of human resource planning: tactical planning and high level (or strategic) planning. On one hand, tactical planning concerns workforce allocation decisions such as within day staffing, allocating capacities among facilities or scheduling operating rooms (Green 2004). A widely studied problem is staff scheduling or “nurse rostering”. The goal is to assign workers to shifts (Felici and Gentile 2004) or tasks (Grunow, Günther and Yang 2004) or to decide which shifts and breaks to use (Aykin 1996) subject to several hard and soft constraints. A summary of solution procedures for the staff scheduling problem appears in Vanhoucke and Maenhout (2007).

Strategic planning, on the other hand, concerns long-term human resource decisions. Those decisions impact the number and type of human resources that will be available over a specified time horizon. They include determining the number of resources to train (Balinsky and Reisman 1972) or cross train (Iravani, Kolfal and Van Oyen 2007) or the number of employees to recruit or promote in either the deterministic setting (Rao 1990) or the stochastic setting (Martel and Price 1981; Chattopadhyay and Gupta 2007) . As discussed by Georgiou and Tsantas (2002) and De Feyter (2007), promotions might occur as a result of a vacancy in a group (pull models) or with a certain probability independent of the current vacancies (push models). In a related application, Wiersma (2007) and Thompson (2007) studied employees’ learning curves and their impact on workforce planning decisions.

This chapter focuses on strategic health human resource planning. Our approach takes into account training, promotion and recruitment decisions. Since nurses represent the largest group of healthcare providers, and were among the top 10 occupations with the largest expected number of employment openings in British Columbia (BC) over this decade (B.C. Stats 2003), we have worked with decision makers to study the policy implications of using high level planning to model the nursing workforce of the province of British Columbia. With relevant data, our model could also be applied to other health professions and jurisdictions.

According to the Canadian Institute for Health Information (2007), some of the key characteristics by which the nursing profession differs from others include:

- A high female to male ratio (which might impact parental leaves).
- Low unemployment rates (as demand for nurses increases, it cannot be assumed that there will be sufficient supply to meet that demand).

- Extreme consequences of having inadequate staff to meet demand (it is not possible to optimize supply without consideration of the demand).
- The presence of regulatory bodies (regulatory bodies may limit entrance of new graduates to the workforce, mobility within Canada and immigration of workers from other countries).
- A skewed age distribution (this might impact expectations of the work environment as well as attrition rates).
- The need for nurses to work both in rural and urban areas (while some professions might provide services remotely, a nurse must be located where the service is needed).
- Data limitations (e.g., standards, privacy, confidentiality).

We describe the innovative features of our model as well as the methodological challenges encountered. We expand models from the operations research literature by incorporating characteristics that are specific to the nursing workforce including full time equivalence (FTE), promotion rules, parental leaves, “on the job” training (continuing education), two types of educational programs and the impact of age on workforce attrition.

## 2.1 Nursing in British Columbia

Canada, like many countries, is experiencing a shortage of registered nurses (RNs) that is projected to worsen over the next decade. This projected shortage is due, in part, to the growth in need for healthcare services by the Canadian population as well as the increased attrition of an aging workforce (Basu and Halliwell 2004). British Columbia’s RNs are the oldest RNs, on average, in the country, and the province has the largest percentage of over-45-year-old nurses (Canadian Institute for Health Information 2006). The large number of RNs expected to retire in the next decade, along with the repercussions of reductions in education seats in the 1990s, have created a significant imbalance between those entering the profession and those leaving (O’Brien-Pallas, Alksnis and Wang 2003; Human Resources and Skills Development Canada 2003). Table 2.1 provides additional information about BC’s population and its RN workforce.



**Table 2.1. BC's employed registered nurses' demographics and population statistics, 1997-2005**

	1997	2001	2005
<b>Demographics of Employed Registered Nurses*</b>	(N = 27,964)	(N = 27,375)	(N = 27,814)
Number per 10,000 Population	70.2	66.7	65.3
Average Age (National Average)	43.3 (42.4)	44.8 (43.7)	46.4 (44.7)
<b>Area of responsibility (Number (%))</b>			
Direct Care	25,723 (92.0)	24,568 (89.7)	24,956 (89.7)
Administration	998 ( 3.6)	1,135 ( 4.1)	1,162 ( 4.2)
Education	1,010 ( 3.6)	1,148 ( 4.2)	1,386 ( 5.0)
Research	146 ( 0.5)	194 ( 0.7)	235 ( 0.8)
Not Stated	87 ( 0.5)	330 ( 1.2)	75 ( 0.3)
<b>Position (Number (%))</b>			
Staff Nurse/Community Health Nurse	22,770 (81.4)	21,819 (79.7)	21,965 (79.0)
Management	2,124 ( 7.6)	2,010 ( 7.3)	2,119 ( 7.6)
Other/Not Stated	3,070 (11.0)	3,546 (13.0)	3,730 (13.4)
<b>Place of Work</b>			
Hospital	18,156 (64.9)	17,599 (64.3)	17,336 (62.3)
Community	3,013 (10.8)	3,273 (12.0)	3,971 (14.3)
Nursing Home	2,333 (8.3)	3,653 (13.0)	3,371 (12.1)
Other/Not Stated	4,462 (16.0)	2,940 (10.7)	3,136 (11.3)
<b>Source of New RNs+ (Number (%))</b>			
BC Schools of Nursing	670 (46.9)	572 (46.7)	857 (55.0)
Other Provinces	601 (42.1)	371 (30.3)	405 (26.0)
Other Countries	158 (11.0)	282 (23.0)	297 (19.0)
<b>BC Population Estimates#</b>			
Total Population	3,948,544	4,078,447	4,257,833
*Canadian Institute for Health Information, based on the number in the registry actively employed in nursing (2002; 2006) + College of Registered Nurses of British Columbia # Statistics Canada			

Estimates based on a cohort analysis conducted by the Canadian Nurses Association (2002) predict that Canada will be short between 78,000 and 113,000 RNs in the next decade. Given this projected deficit, strategies addressed specifically to those entering the profession and those who will be retiring from the profession are of particular importance. Increasing education seats is considered a key strategy to address the shortage; "governments must be engaged immediately with schools and employers, educating health professionals to put in place the human and physical resources to accommodate more students" (Villeneuve and MacDonald 2006:101). Yet the number of students that should or could be accommodated is an open question.

Education seats are not the only consideration in planning human resources. Multiple factors influence the supply of, and demand for, RNs, including attrition (short-term leaves, premature leave from the profession and retirement), changes in the workplace (e.g., availability of employment positions and contractual requirements such as hours worked), the availability of personnel in leadership or management roles, the skill mix and task delegation of all healthcare providers, how demand is defined (whether based on current utilization patterns or population healthcare needs) and productivity (availability of support staff and other aids).

Governmental and employers' policies also play an important role in planning human resources. The need to examine how policy decisions made in one sector of the healthcare system can influence other sectors is an issue that is rarely considered. For example, a decision to increase education seats should be commensurate with the availability of employment positions. Studies by the College of Registered Nurses of British Columbia (2006) have indicated that the availability of full-time employment is a factor in whether recently graduated RNs remain in the province.

Recruitment and retention strategies are an important feature of planning human resources. A growing body of literature has identified the importance of workplace characteristics for RN retention. In their synthesis of research and other literature on nursing and work, Baumann et al. (2001) identified features of the work environment related to perceived quality, including the availability of personnel in leadership or management positions and an emphasis on promoting recruitment and retention. Examining the influence of various strategies on the overall supply of RNs is an important element of planning (Kephart et al. 2006).

## **2.2 The Model**

The goal of the model described here was to investigate the impact of assumptions on the number of students to admit to all university and college programs in the province, the number of nurses to train for management roles, and the number of nurses to recruit from outside the province to meet forecasted service needs. The model variables are totals and do not represent specific individuals so that constraints involving proportions are equivalent to taking expected values or averaging over a large population.

The problem was formulated as a multi-period finite horizon linear program (LP). LP was chosen over other approaches because of its transparency, the capability of obtaining an optimum solution, as well as its simplicity to modify in order to perform "what-if" analyses. Alternatives considered

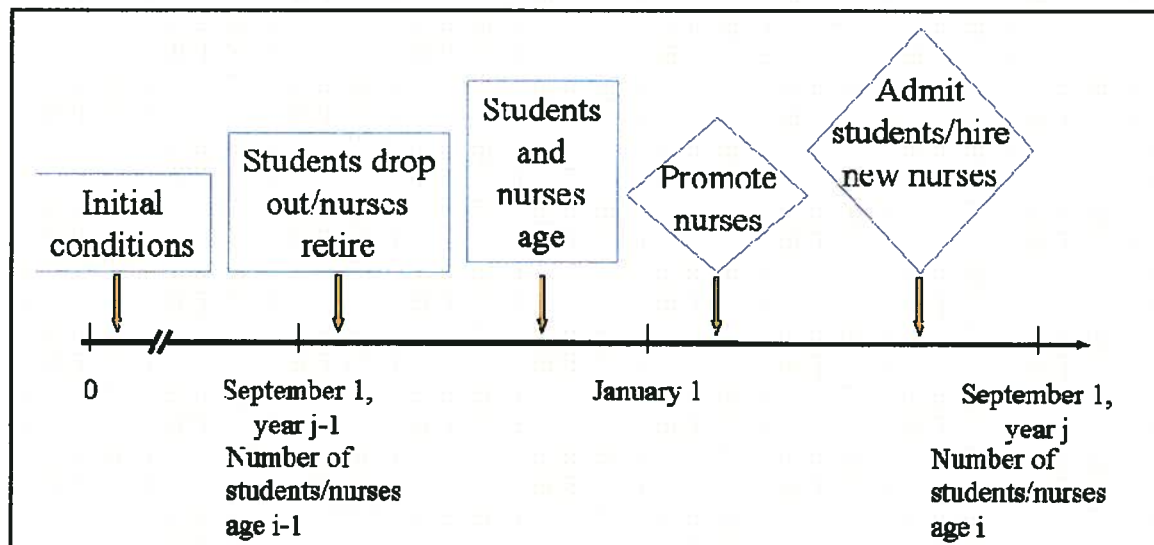
include simulation at the cohort level, simulation at the individual level and stochastic programming.

### 2.2.1 Model Assumptions

The model was based on the year to year dynamics of the entire nursing cohort broken down by age and activity level. Consequently the model was based on mathematical expectations, which are equivalent to averaging over a large population. While we acknowledge that this is a strong assumption, it is justified by the law of large numbers and appropriate for aggregate level work force policy decisions. The effect of variability on our optimal solutions could be investigated using simulation or stochastic programming.

The timeline of events and decisions in the model appears in Figure 2.1. It represents the following chain of events.

**Figure 2.1. Model timeline**

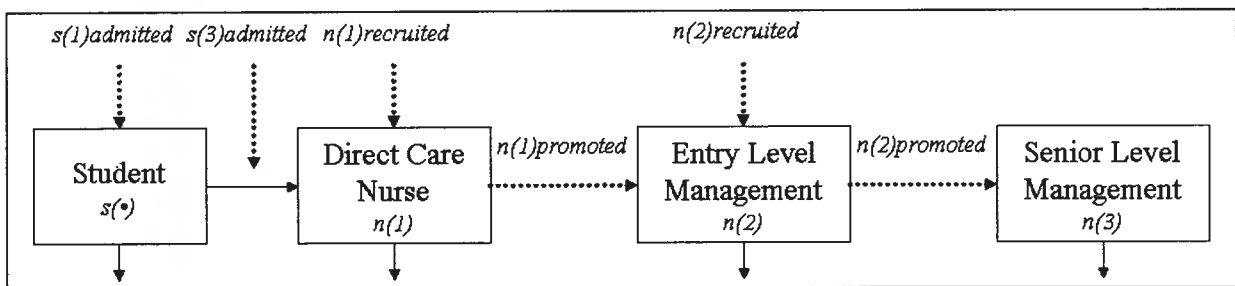


An academic year starts with a known number of students and nurses in each age category and at each level. For students, the level represents their year in their educational program, and for nurses, it represents whether they are practicing or whether they are employed at one of two managerial levels. Throughout the year, students and nurses leave the system through attrition. Further the student and nursing populations age, and students who do not drop out advance in their educational programs or graduate. Then, a decision is made regarding how many new students to admit, how many nurses to promote to managerial positions, and how many nurses to

recruit from outside the province. The results of this decision, together with the sizes of the age cohorts (at January 1 in the figure), determines the student and nursing populations in the following year.

A system flowchart appears in Figure 2.2. Dashed arrows represent decision variables. The model incorporates the two types of baccalaureate educational programs offered in the province: the standard program and the advanced standing program. Students that enrol in the standard program take four years to complete their education; they are not required to have any post-secondary education prior to being admitted into a program. Students admitted to the advanced standing program have completed a minimum number of university credits, *a priori*. Those students are assumed to be older than the students admitted to the standard program. The advanced standing program is two years in length, and its content is similar to the last two years of the standard program.

**Figure 2.2. Workforce flowchart**



Some students discontinue their education (attrition being highly dependent on the year of study (Pringle and Green 2005)), while most students graduate with a nursing degree. A small fraction of graduates might decide to leave the province or not register with the regulatory body. To practice as a registered nurse in the province, graduates of a nursing program must pass the Canadian Registered Nurse Examination (College of Registered Nurses of British Columbia 2006). Direct care nurses that have acquired the relevant experience might be promoted to work in entry level management positions and subsequently to senior level management positions. Nurses might leave the profession permanently or temporarily (parental leaves being a key factor included in the model); attrition rates are highly dependent on age (Kazanjan et al. 1986).

In addition to graduates of BC schools of nursing, other sources of new RNs in British Columbia are other provinces and countries. We assume that this migration, which is affected by the global

shortage of nurses, will continue in the near future. High recruitment costs (Webber 2005) as well as availability of nurses willing to move to the province are incorporated into the model.

### 2.2.2 Decision Variables

The decision variables of the model are the total number of students to admit into the standard and advanced standing programs ( $s(1)admitted_j$  and  $s(3)admitted_j$  respectively), the total number of direct care nurses and entry level managers recruited ( $n(1)recruited_j$  and  $n(2)recruited_j$  respectively), the total number of direct care nurses promoted to entry level management positions ( $n(1)promoted_j$ ), and the total number of entry level managers promoted to senior level management positions ( $n(2)promoted_j$ ). Given that certain characteristics are age dependent (such as attrition rates and fertility rates), it is necessary to keep track of the number of first, second, third and fourth year students of age  $i$  ( $A \leq i \leq B$ ) in year  $j$  ( $s(1)_{ij}$ ,  $s(2)_{ij}$ ,  $s(3)_{ij}$ ,  $s(4)_{ij}$ ) as well as the total number ( $n(1)_{ij}$ ,  $n(2)_{ij}$ ,  $n(3)_{ij}$ ) of direct care nurses, entry level managers and senior level managers of age  $i$  in year  $j$ . We call these variables “bookkeeping variables” and calculate them as follows:

The total number of first year students of age  $i$  at time period  $j$  ( $s(1)_{ij}$ ) is given by:

$$s(1)_{ij} = s(1)admitted_j * Ps(1)_i \quad \forall i, j \quad (2.1)$$

where  $Ps(1)_i$  represents the probability that a student admitted into the first year of the program is of age  $i$ . Note that we are assuming that the age distribution of the students admitted into the program remains the same over time, that is, it does not depend on the period  $j$ .

The total number of second year students of age  $i$  at time period  $j$  ( $s(2)_{ij}$ ) is given by:

$$s(2)_{ij} = initial\_s(2)_i \quad \text{for } j = 1, A \leq i \quad (2.2)$$

$$s(2)_{ij} = s(1)_{i-1, j-1} * Pcontinuing(1) \quad \text{for } 2 \leq j \leq N, A \leq i < B \quad (2.3)$$

$$s(2)_{ij} = [s(1)_{i-1, j-1} + s(1)_{ij-1}] * Pcontinuing(1) \quad \text{for } 2 \leq j \leq N, i = B \quad (2.4)$$

While equation (2.2) specifies that the number of second year students in the first period is just based on the initial conditions  $initial\_s(2)_i$  (e.g., decisions made in the first period will only affect the number of second year students after the second period), equations (2.3) and (2.4) indicate that the number of second year students is a fraction of the number of students that were admitted

into the standard program the year before. Equation (2.4) can be interpreted by noting that, after students reach age B, their attrition rate does not change as they age further.

The total number of third year students of age  $i$  at time period  $j$  ( $s(3)_{ij}$ ) is given by:

$$s(3)_{ij} = \text{initial\_}s(3)_i + s(3)\text{admitted}_j * Ps(3)_i \quad \text{for } j = 1, A \leq i \quad (2.5)$$

$$s(3)_{ij} = s(2)_{i-1,j-1} * P\text{continuing}(2) + s(3)\text{admitted}_j * Ps(3)_i \quad \text{for } 2 \leq j \leq N, A \leq i < B \quad (2.6)$$

$$s(3)_{ij} = [s(2)_{i-1,j-1} + s(2)_{ij-1}] * P\text{continuing}(2) + s(3)\text{admitted}_j * Ps(3)_i \quad \text{for } 2 \leq j \leq N, i = B \quad (2.7)$$

Equations (2.5) through (2.7) compute the total number of third year students by summing the number of students that are in the standard program who continue after their second year, and the number of students admitted into the advanced standing program.

Finally, the total number of fourth year students of age  $i$  at time period  $j$  ( $s(4)_{ij}$ ) is given by:

$$s(4)_{ij} = \text{initial\_}s(4)_i \quad \text{for } j = 1, A \leq i \quad (2.8)$$

$$s(4)_{ij} = s(3)_{i-1,j-1} * P\text{continuing}(3) \quad \text{for } 2 \leq j \leq N, A \leq i < B \quad (2.9)$$

$$s(4)_{ij} = [s(3)_{i-1,j-1} + s(3)_{ij-1}] * P\text{continuing}(3) \quad \text{for } 2 \leq j \leq N, i = B \quad (2.10)$$

Equations (2.8) through (2.10) can be interpreted in the same way as equations (2.2) though (2.4).

The total number of direct care nurses of age  $i$  at time period  $j$  ( $n(1)_{ij}$ ) is given by:

$$n(1)_{ij} = \text{initial\_}n(1)_i + n(1)\text{recruited}_j * Prn(1)_i - n(1)\text{promoted}_j * Ppn(1)_i \quad \text{for } j = 1, A \leq i \leq B \quad (2.11)$$

$$n(1)_{ij} = n(1)_{i-1,j-1} * (1 - P_{\text{retire}}(1)_{i-1}) + s(4)_{i-1,j-1} * P\text{continuing}(4) * P_{\text{pass}} * P_{\text{stay}} + n(1)\text{recruited}_j * Prn(1)_i - n(1)\text{promoted}_j * Ppn(1)_i \quad \text{for } 2 \leq j \leq N, A \leq i < B \quad (2.12)$$

$$\begin{aligned}
n(1)_{ij} = & n(1)_{i-1,j-1} * (1 - P_{retire}(1)_{i-1}) \\
& + n(1)_{i,j-1} * (1 - P_{retire}(1)_i) \\
& + s(4)_{i-1,j-1} * P_{continuing}(4) * P_{pass} * P_{stay} \\
& + s(4)_{i,j-1} * P_{continuing}(4) * P_{pass} * P_{stay} \\
& + n(1)_{recruited_j} * Prn(1)_i \\
& - n(1)_{promoted_j} * Ppn(1)_i \quad \text{for } 2 \leq j \leq N, i = B
\end{aligned} \tag{2.13}$$

In the first period,  $n(1)_{ij}$  is calculated by subtracting the direct care nurses that are promoted from the initial number of direct care nurses  $initial\_n(1)_i$  and then adding recruits by age (equation (2.11)). In the remaining periods (equations (2.12) and (2.13)), the total number of direct care nurses is comprised of those nurses that are neither promoted nor decide to retire, plus new recruits and students that complete their education ( $P_{continuing}(4)$ ), pass the national examination ( $P_{pass}$ ), and practice in the province ( $P_{stay}$ ). We assumed that the age distribution of nurses recruited  $Prn(1)_i$  and promoted  $Ppn(1)_i$  was known and constant over the planning horizon. Moreover, after nurses reach age B, their attrition rate does not change as they age further.

The total entry level managers of age  $i$  at time period  $j$  ( $n(2)_{ij}$ ) is given by:

$$\begin{aligned}
n(2)_{ij} = & initial\_n(2)_i + n(1)_{promoted_j} * Ppn(1)_i \\
& + n(2)_{recruited_j} * Prn(2)_i \\
& - n(2)_{promoted_j} * Ppn(2)_i \quad \text{for } j = 1, A \leq i \leq B
\end{aligned} \tag{2.14}$$

$$\begin{aligned}
n(2)_{ij} = & n(2)_{i-1,j-1} * (1 - P_{retire}(2)_{i-1}) \\
& + n(1)_{promoted_j} * Ppn(1)_i \\
& + n(2)_{recruited_j} * Prn(2)_i \\
& - n(2)_{promoted_j} * Ppn(2)_i \quad \text{for } 2 \leq j \leq N, A \leq i < B
\end{aligned} \tag{2.15}$$

$$\begin{aligned}
n(2)_{ij} = & n(2)_{i-1,j-1} * (1 - P_{retire}(2)_{i-1}) \\
& + n(2)_{i,j-1} * (1 - P_{retire}(2)_i) \\
& + n(1)_{promoted_j} * Ppn(1)_i \\
& + [n(2)_{recruited_j} - n(2)_{promoted_j}] * Pn(2)_i \quad \text{for } 2 \leq j \leq N, i = B
\end{aligned} \tag{2.16}$$

The interpretation of equations (2.14) through (2.16) is similar to the interpretation of equations (2.11) through (2.13). The main difference is that, while new direct care nurses from within the province are a function of the students from the previous period, new entry level managers from within the province are based on the number of direct care nurses promoted.

Lastly, the total senior level managers of age  $i$  at time period  $j$  ( $n(3)_{ij}$ ) is given by:

$$n(3)_{ij} = \text{initial\_}n(3)_i + n(2)\text{promoted}_j * Pp_n(2) \quad \text{for } j = 1, A \leq i \leq B \quad (2.17)$$

$$n(3)_{ij} = n(3)_{i-1,j-1} * (1 - P_{\text{retire}}(3)_{i-1}) + n(2)\text{promoted}_j * Pn(2)_i \quad \text{for } 2 \leq j \leq N, A \leq i < B \quad (2.18)$$

$$n(3)_{ij} = n(3)_{i-1,j-1} * (1 - P_{\text{retire}}(3)_{i-1}) + n(3)_{ij-1} * (1 - P_{\text{retire}}(3)_i) + n(2)\text{promoted}_j * Pn(2)_i \quad \text{for } 2 \leq j \leq N, i = B \quad (2.19)$$

Note that we assumed that it was not possible to recruit senior level managers: all senior level managers must have worked as entry level managers (equations (2.17) through (2.19)).

To determine the effective workforce size each year, our model adjusted the workforce size as follows. The number of full time equivalents each year,  $FTE(n(\bullet)_{ij})$ , was a function of the total number of direct care nurses, entry level managers and senior level managers. For instance, to incorporate parental leaves, we subtracted the expected number of workers of each age that requested such leave in a given year times the fraction of the year the nurse would be away if such leave was granted. Given the high ratio of females in the occupation (over 94% (Canadian Institute for Health Information 2007)), full time equivalence due to parental leave  $FTE(n(\bullet)_{ij})_{pl}$  could be approximated by:

$$FTE(n(\bullet)_{ij})_{pl} = n(\bullet)_{ij} * (1 - \text{Fertility\_rate}_i * \text{female\_ratio} * \text{Maternity\_leave}/12) \quad (2.20)$$

where  $\text{Fertility\_rate}_i$  represents the ratio of females of age  $i$  that request parental leave per year,  $\text{female\_ratio}$  is the female to male ratio and  $\text{Maternity\_leave}$  is the number of months of leave specified by labour contract.

In addition, we assumed that nurses that started a new position would take some time to get adjusted to their position (due to the learning curve and the time required to take the national examination). Hence, the full time equivalents  $FTE()$  of those nurses was computed as a fraction of the nurses that started the new position.

### 2.2.3 Data Formulation

A summary of the data required to populate the model appears in Table 2.2.



Table 2.2. Notation of inputs

	Notation	Description	Source
Probabilities	$Ps(1)_i, Ps(3)_i$	Probability that a student admitted into the program is of age $i$	(University of British Columbia, School of Nursing 2006)
	$Prn(1)_i, Prn(2)_i$	Age distribution of nurses recruited	(College of Registered Nurses of British Columbia 2006) ; (Canadian Institute for Health Information 2006)
	$Ppn(1)_i, Ppn(2)_i$	Age distribution of nurses promoted	(College of Registered Nurses of British Columbia 2006) ; (Canadian Institute for Health Information 2006)
	$Pcontinuing(1),$ $Pcontinuing(2),$ $Pcontinuing(3),$ $Pcontinuing(4)$	Fraction of students that continue in the program each year	UBC School of Nursing; (Pringle and Green 2005)
	$Ppass$	Fraction of BC graduates that pass the national CRNE examination	(College of Registered Nurses of British Columbia 2006)
	$Pstay$	Fraction of RNs that remain in the province after graduation	(College of Registered Nurses of British Columbia 2006)
	$P_{retire}(1)_i,$ $P_{retire}(2)_i,$ $P_{retire}(3)_i$	Attrition rate of nurses by age $i$	(Kazanjan, Brothers and Wong 1986); (O'Brien-Pallas, Alksnis and Wang 2003)
	$Fertility\_rate_i$	Ratio of females of age $i$ that request parental leave per year	(Statistics Canada 2004)

	Notation	Description	Source
Initial conditions	$initial\_s(2)_i$ $initial\_s(3)_i$ $initial\_s(4)_i$	Number of students of age $i$ in the first period of the model	UBC School of Nursing
	$initial\_n(1)_i$ $initial\_n(2)_i$ $initial\_n(3)_i$	Number of nurses of age $i$ in the first period of the model	(College of Registered Nurses of British Columbia 2006) ; (Canadian Institute for Health Information 2006)
	$pfraction(1)_x$ $pfraction(2)_x$ $pfraction(3)_x$	Initial fraction of workers that have been in their position at least $x$ years	(College of Registered Nurses of British Columbia 2006) ; (Canadian Institute for Health Information 2006)
	$female\_ratio$	Female to male ratio	(Canadian Institute for Health Information 2006)
Costs	$tsCost$	Cost of funding a university seat per year	BC Ministry of Advanced Education (MAE)
	$tn(2)Cost$ $tn(3)Cost$	Cost of training to promote a nurse from direct care RN into a managerial position	MAE
	$rn(1)Cost_j$ $rn(2)Cost_j$	Recruitment cost for each nurse in year $j$	(Webber 2005)
	$sn(1)Cost$ $sn(2)Cost$ $sn(3)Cost$	Average annual salaries	BC Nurses' Union
Bounds	$min\ s(1)_i$ $min\ s(3)_j$	Lower bound on the number of students admitted into the programs in year $j$	UBC School of Nursing

	Notation	Description	Source
Bounds	$max\ s(1)_j$ $max\ s(3)_j$	Upper bound on the number of students admitted into the programs in year $j$	UBC School of Nursing
Demand	$BCpop_j$	Population projection for year $j$	(BC Stats 2006)
	$n(1)ratio_j$ , $n(2)ratio_j$ , $n(3)ratio_j$	Minimum ratios of nurses to meet population demand in year $j$	(Canadian Institute for Health Information 2006)

Obtaining this data was challenging and time consuming. For this purpose, we performed an extensive literature review, and as well contacted various organizations including the British Columbia Nurses' Union, the College of Registered Nurses of British Columbia, the BC Ministry of Advanced Education, and the School of Nursing at the University of British Columbia. Since all nurses in BC are registered to practice, the authors investigated the movement of registrants between membership statuses as a proxy for employment status.

Other key inputs to the model were the nurse/population and the managerial ratios. We first set as our objective to maintain current ratios (as estimated by the Canadian Institute for Health Information (2006)) and performed a sensitivity analysis to determine the cost of improving such ratios. However, it must be highlighted that the ratios were only used as proxies for demand. Given the key role they play in the solution, a more refined needs-based model should be developed.

#### 2.2.4 Model Formulation

In this subsection, we first show the full model and then give further detail on how the objective function and constraints were calculated. The mathematical formulation of the model is as follows:

$$\text{Min } \sum_j \text{Cost of Training}_j + \text{Recruitment Cost}_j + \text{Annual Salary}_j$$

subject to:

Necessary resources:

$$max\ s(1)_j \geq s(1)admitted_j \geq min\ s(1)_j \quad \forall j \quad (2.21)$$

$$max\ s(3)_j \geq s(3)admitted_j \geq min\ s(3)_j \quad \forall j \quad (2.22)$$

$$\sum_i FTE(n(1)_{i,j}) \geq BCpop_j / n(1)ratio_j \quad \forall j \quad (2.23)$$

$$\sum_i FTE(n(2)_{i,j}) \geq \sum_i FTE(n(1)_{i,j}) / n(2)ratio_j \quad \forall j \quad (2.24)$$

$$\sum_i FTE(n(3)_{i,j}) \geq \sum_i FTE(n(1)_{i,j}) / n(3)ratio_j \quad \forall j \quad (2.25)$$

Balance first and third year students:

$$s(1)admitted_j \geq s(3)admitted_j \quad \forall j \quad (2.26)$$

Do not recruit more direct care nurses than the number of graduates that can work in BC:

$$\sum_i initial\_s(4)_i * Pcontinuing(4) * Ppass * Pstay \geq n(1)recruited_j \quad j = 1 \quad (2.27)$$

$$\sum_i s(4)_{i,j-1} * Pcontinuing(4) * Ppass * Pstay \geq n(1)recruited_j \quad \forall j \geq 2 \quad (2.28)$$

Only promote those nurses that have been in their position x years:

$$n(1)promoted_j \leq pfraction(1)_x * \sum_i n(1)_{i,0} \quad j = 1 \quad (2.29)$$

$$n(1)promoted_j \leq \sum_i \left[ n(1)_{i,0} * pfraction(1)_{x-j} * \prod_{k=0}^{j-2} (1 - P\_retire(1)_{i+k}) \right] \quad x \geq j \geq 2 \quad (2.30)$$

$$n(1)promoted_j \leq \sum_i \left[ n(1)_{i,j-x} * \prod_{k=0}^{x-1} (1 - P\_retire(1)_{i+k}) \right] \quad j > x \quad (2.31)$$

$$n(2)promoted_j \leq pfraction(2)_x * \sum_i n(2)_{i,0} \quad j = 1 \quad (2.32)$$

$$n(2)promoted_j \leq \sum_i \left[ n(2)_{i,0} * pfraction(2)_{x-j} * \prod_{k=0}^{j-2} (1 - P\_retire(2)_{i+k}) \right] \quad x \geq j \geq 2 \quad (2.33)$$

$$n(2)promoted_j \leq \sum_i \left[ n(2)_{i,j-x} * \prod_{k=0}^{x-1} (1 - P\_retire(2)_{i+k}) \right] \quad j > x \quad (2.34)$$

The objective of the model is to minimize the total cost of training, recruitment and annual salaries. The cost of training is comprised of the cost of educating each student (e.g.,  $tsCost$ ) and the cost of promoting nurses into managerial positions (e.g.,  $tn(2)Cost$  and  $tn(3)Cost$ ). The cost of training in a given period  $j$  can therefore be written as:

$$\begin{aligned} Cost\ of\ Training_j = & \sum_i [s(1)_{i,j} + s(2)_{i,j} + s(3)_{i,j} + s(4)_{i,j}] * tsCost \\ & + n(1)promoted_j * tn(2)Cost + n(2)promoted_j * tn(3)Cost \end{aligned} \quad (2.35)$$

In order to calculate the annual recruitment cost, we multiplied the expected number of health professionals recruited in each position by the recruitment cost ( $rn(1)Cost_j$  for direct care nurses and  $rn(2)Cost_j$  for entry level managers):

$$Recruitment\ Cost_j = n(1)recruited_j * rn(1)Cost_j + n(2)recruited_j * rn(2)Cost_j \quad (2.36)$$

Finally, annual salaries were calculated by summing the estimates of annual salaries for the various positions (e.g.  $sn(1)Cost$ ,  $sn(2)Cost$  and  $sn(3)Cost$ ):

$$Annual\ Salary_j = \sum_i n(1)_{i,j} * sn(1)Cost + \sum_i n(2)_{i,j} * sn(2)Cost + \sum_i n(3)_{i,j} * sn(3)Cost \quad (2.37)$$

We now interpret the above constraints:

Necessary resources constraints: These constraints correspond to the bounds on the expected number of students and nurses in the province. Constraints (2.21) and (2.22) indicate that the expected number of students admitted into the standard and advanced standing educational programs should be within an upper and lower bound ( $max\ s(\bullet)_j$  and  $min\ s(\bullet)_j$ ). These bounds could be absolute numbers or a function of the students admitted the year(s) before. The constraints represent various challenges faced when expanding/decreasing the university programs (which include the need for faculty and infrastructure changes).

Constraints (2.23), (2.24) and (2.25) drive the model. Constraint (2.23) ensures that there are a sufficient number of full time equivalent direct care nurses to meet annual targets over the planning horizon. While we currently estimate the minimum number of nurses to meet the population's need in a given year based on population projections, demand could be defined by rates of hospital utilization or other measures of population healthcare needs. Moreover, constraints (2.24) and (2.25) define minimum nurse to manager ratios. It is widely acknowledged that nurse to manager ratios play a key role in nurse retention (Kramer, Schmalenberg and Maguire 2004). Our model is based on current ratios and attritions. Sensitivity analysis was used to investigate the impact of altering such ratios.

Balance first- and third-year student constraints: Students that are admitted into the advanced standing educational program must previously have completed a minimum number of university credits. Given this requirement, and the assumption that such education must be received within

the province, current practice is not to admit more students to the advanced standing program than into the standard program. Constraint (2.26) ensures that this condition is satisfied. This constraint could easily be relaxed to reflect policy changes.

Do not recruit more direct care nurses than the number of graduates that can work in BC constraint: As was previously mentioned, the number of nurses that can be recruited is constrained by the availability of nurses to come and work in the province. Historically, over 50% of newly employed nurses are graduates of BC. Constraints (2.27) and (2.28) ensure that this ratio continues in the future. Such constraints could be modified to address the impact that the expected global shortage of health professionals might have on the entire system.

Only promote those nurses that have been in their position over x years constraint: Finally, constraints (2.29) through (2.34) ensure that nurses acquire the experience necessary before being promoted to a managerial position. While constraints (2.31) and (2.34) indicate that the number of nurses that can be promoted in period  $j$  is a function of the number of nurses that were working  $x$  periods before and that did not retire (attrition probability by age is expressed as  $Pretire(\bullet)_i$ ), constraints (2.29), (2.30), (2.32) and (2.33) require knowledge of the initial fraction of resources that have been in their position at least  $x$  years ( $pfraction(\bullet)_x$ ).

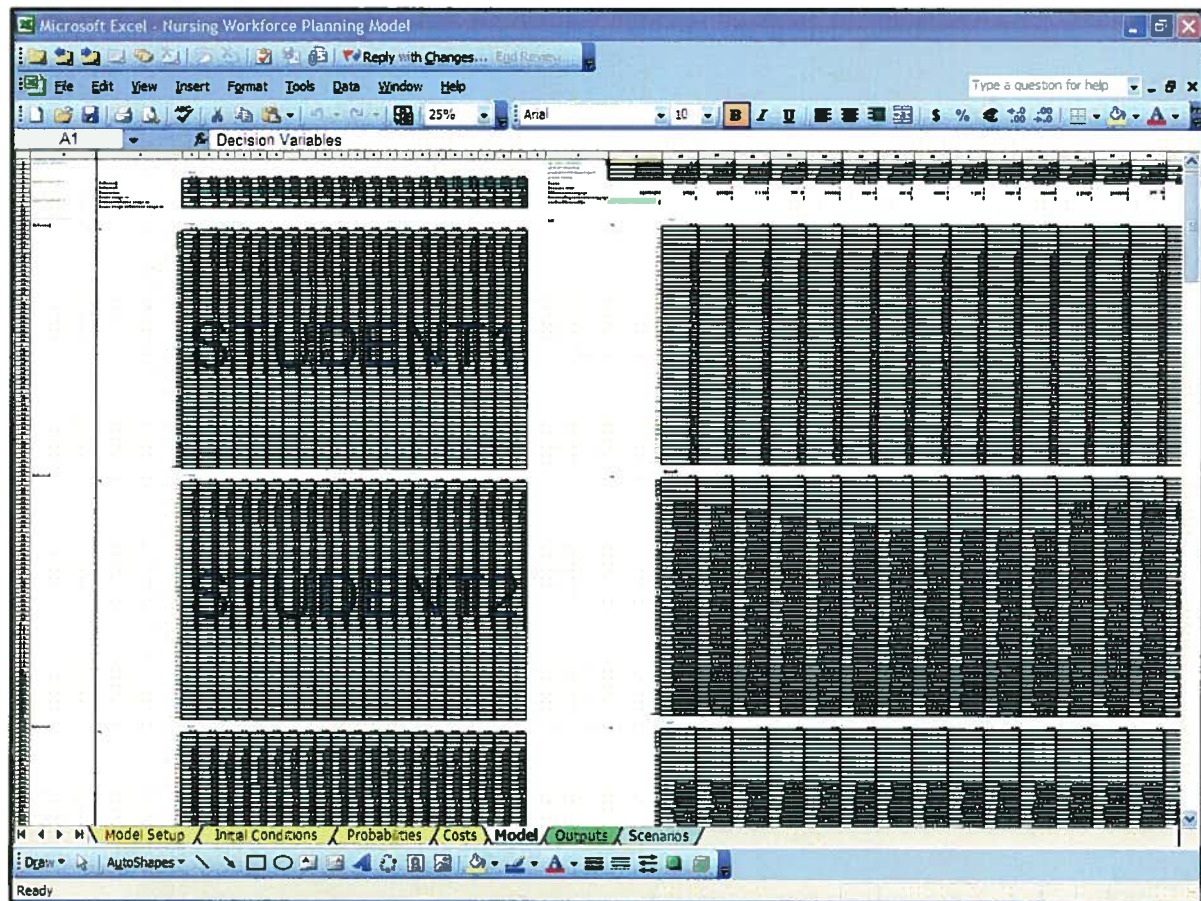
We note also that an impact of using this objective function was that solutions were to satisfy lower bound constraints if feasible. Thus we obtained solutions that used the smallest workforces possible.

## 2.3 Implementation

This model was applied to the BC registered nurse workforce. As was previously mentioned, a key challenge was encountered in obtaining the necessary data to populate the model. The main inputs to the model were the objective function coefficients (cost of training, recruitment cost and annual salaries) as well as the constraints coefficients (age demographics, attrition rates, workforce ratios, demand for providers and full time equivalence basis). The sources of data for our model are summarized in Table 2.2. They included government, education programs, regulatory bodies, census data, employers (including private sector), research literature and unions and professional associations.

A user friendly interface was designed in the MS-Excel® platform, and the Frontline solver add-on (Frontline Systems, Inc. 2007) was used to obtain the solutions. A screen shot of the model appears in Figure 2.3. The initial model was solved over 20 periods with 50 age categories. This model had 126 variables and 7497 bookkeeping variables. We were able to obtain an optimal solution (or determine that the model was infeasible) within seconds on a PC.

**Figure 2.3. Model screenshot**



### 2.3.1 Scenario Analysis

While the model was used to determine the optimal education, promotion and recruitment plan and its associated costs, the advantage of this approach relies on its ability to perform “what-if?” analyses. This was particularly useful when faced with challenges in the data availability and reliability. Five scenarios are analyzed herein. They were chosen to address current policy concerns and to illustrate how the model could be used for setting policy. In each case, we made the appropriate changes to the model inputs and solved it to find an optimal education,

recruitment and promotion policy. We then compared the model output graphically. We used the model to determine which factors had the largest impact on the long term workforce plan.

### Scenario 1 – Baseline

The first scenario assumes:

- The goal is to satisfy current provider-to-population ratios subject to increasing populations.
- In the first 10 years of the model, the percentage increase in the number of students admitted to university programs is constrained to be the same as the maximum increase that has been observed in the past.
- After entering the workforce, RNs who have completed the advanced standing program have the same attrition rates, by age, as RNs that have completed a standard program.
- When RNs are initially employed they require time for adjustment to the position (i.e., in their first year of a new position they work 0.8 FTEs to accommodate orientation, mentorship and learning).
- A minimum of 500 direct care RNs move to BC each year.
- There are no restrictions on the maximum number of RNs that can be recruited from other provinces and countries in the first year. However, we impose an increased cost associated with bringing new direct care RNs to the province.

### Scenario 2 – Changes in educational program attrition rates

The second scenario addresses the impact of changes in the proportion of students who continue in their educational program after each year of study. The attrition rates are the ratio of students in two consecutive years of schooling.

Reported attrition rates from nursing educational programs vary widely and have been noted to range between 3% and 44% (Pringle and Green 2005). We tested a range of possible attrition rate scenarios. At baseline (Scenario 1), attrition rates of 10% in the first year and 2% and 5% in the second and third years, respectively, were used. In this scenario, we investigated how the optimal policy would change if there were no attrition. Such a scenario might not be realistic, but it shows the impact of reducing educational program attrition rates on other decision variables.



### Scenario 3 – Simultaneously change the direct care RN-to-manager ratio and the practicing RN attrition rate

The baseline scenario assumes a direct care RN-to-manager ratio that follows the national average of approximately 50 direct care RNs and 4 entry-level managers per senior-level manager (Canadian Institute for Health Information 2006). We chose to investigate the impact of two simultaneous changes to the baseline scenario: reducing the direct care RN-to-manager ratio by 10% while assuming that the change would reduce the attrition rates of all RNs by 10%. We caution the reader that the effect of reducing RN to manager ratios on attrition rates has not been widely studied. Although many researchers acknowledge the importance of managers for direct care RN retention (Kramer et al. 2004), they do not report the ideal ratio. Therefore, empirical research is needed to determine if reductions in RN attrition could be achieved by altering direct care RN to manager ratios.

### Scenario 4 – Change the length of parental leave

The length of parental leave entitlement is a controversial topic. Although we do not advocate for longer or shorter parental leaves, we used the model to investigate the impact on recruitment, promotion and training of shortening parental leaves to six months from the 12 months assumed in the baseline scenario. We assumed that the annual fertility rate of RNs was the same as the population average fertility rate, by age, of all women in British Columbia in 2004 (Statistics Canada 2004). This approach fails to recognize the rate at which male RNs and adoptive parents exercise their parental leave provisions.

### Scenario 5 – Change the RN-to-population ratio

In the absence of demand variables for the model, we used the RN to population ratio as a proxy for demand. There is no consensus, however, on what the “proper” ratio ought to be. As mentioned previously, the model allows for demand to be defined by the decision maker. For example, if demand is defined by rates of hospital utilization or population healthcare needs, then these variables could be entered. We therefore believe that the model is a complement to a demand-based model that provides more informed estimations of the minimum number of RNs needed to meet the population’s health needs in a given year.

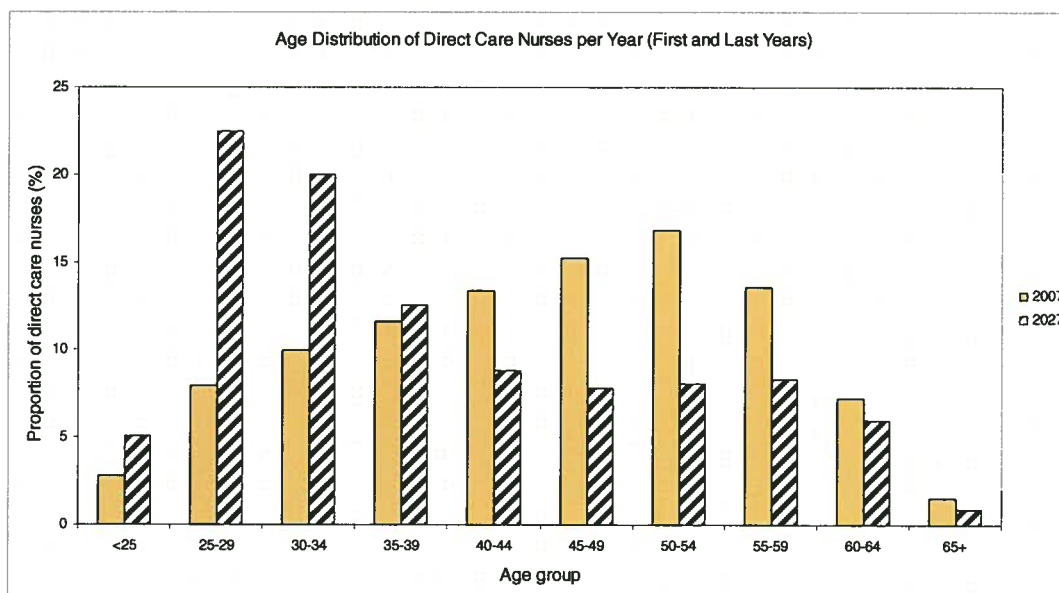
Suppose that a needs-based model suggested that a reasonable target was to have five fewer people per RN than the baseline, and that in such a case, the attrition rate of RNs was reduced by 15%. Under these assumptions, not only would the total number of direct care RNs increase in the long term, but it could be done with a lower yearly recruitment rate. Other possible scenarios include a change in the ratio over the planning horizon (which might be associated with changes in the age distribution of the population), or the direct input of the minimum number of RNs needed based on the population's health status or other characteristics.

### 2.3.2 Results

Using the model, we showed that current policy in British Columbia is not sustainable over the long run. Given our initial conditions, only by having a large initial recruitment or by sacrificing the quality of service (e.g., by changing nurse-to-population ratios) were we able to obtain a feasible solution.

The initial workforce age distribution had a significant impact on the model solution. This is due to the increased attrition of the aging nursing workforce (Basu and Halliwell 2004). Figure 2.4 shows that the age distribution changes to meet long term needs. A next step in this approach would be to model the best sustainable age distribution at each level.


**Figure 2.4. Age distribution of direct care nurses in the first and last years of the planning horizon based on the optimal solution to the model**



Optimal values for the policy variables under the five scenarios are represented in Figures 2.5 through 2.8.

Figure 2.5 shows the total number of direct care RNs available each year in BC under each of the scenarios. The minimum number of RNs required to maintain the current RN to population ratio has also been included. Note that in none of the scenarios does the model suggest having exactly the targeted number of RNs in the first years. This is because planning decisions this year will have long-term implications and consequently impact planning decisions in the future. Note also that in all of the scenarios, the rate of increase in the total number of direct care RNs stabilizes after a few years. Also, as the current workforce is older than is optimal, adding RNs early on is necessary to meet future needs for RNs and managers.

**Figure 2.5. Number of direct care registered nurses per year**

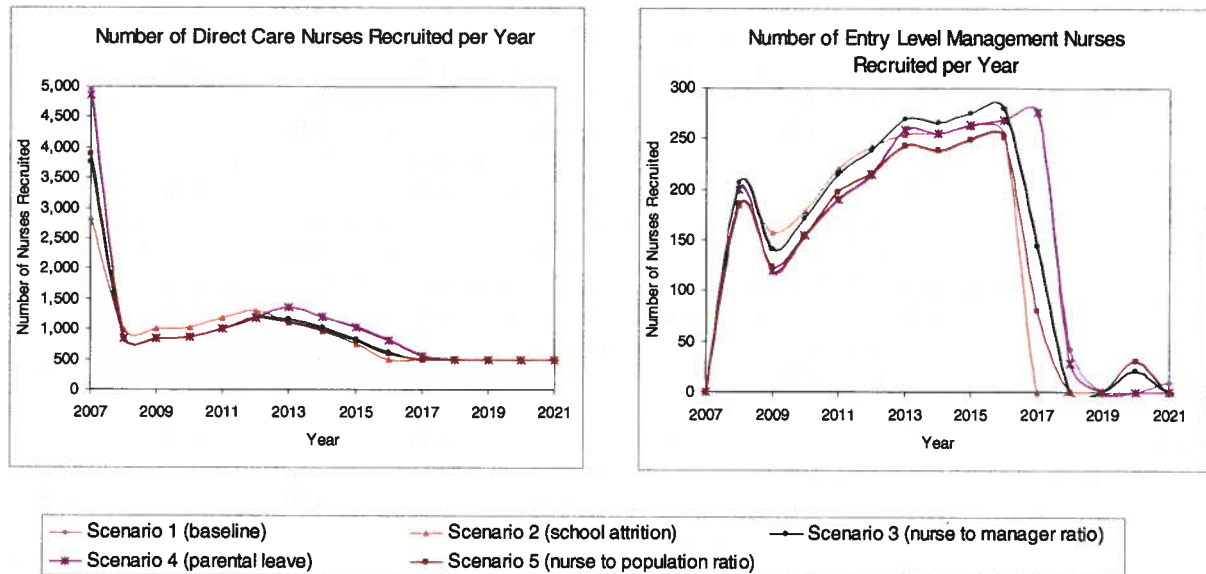


Even with an increase of entry-level managers and senior-level managers (Scenario 3), the total number of new RNs that would be needed is lower than if such an increase were accompanied by a reduction in attrition from the profession. Furthermore, note that the reduction in duration of parental leave did not have a significant effect on the solution. We also tested a scenario with no parental leave (an extreme scenario acknowledged not to be feasible) and noticed no major difference in the number of students admitted to educational programs or recruitment requirements. This may be related to the low fertility rates in the province. However, if it were true that radically reducing parental leave would not reduce the need to educate or recruit new RNs, and if by having such a benefit, attrition rates for RNs could be reduced, Scenario 4 supports a decision to increase parental leave from 6 to 12 months as was recently done in Canada. We

emphasize that to make a decision based on this scenario, other factors involving parental leave (such as the possible association with the fertility rate) would need to be considered.

Figure 2.6 summarizes yearly recruitment requirements. Observe that a large recruitment of direct care RNs from outside the province is required in all of the scenarios, especially in the first year. While the number of direct care RNs recruited is much lower after the first year, note that the recruitment number of direct care RNs is still elevated up to the year 2016. This occurs because the current workforce is not sufficient to meet short term needs given the current RN age distribution and attrition rates. Although we assume that a higher cost is associated with the recruitment of RNs from outside the province compared with the education of new RNs within the province, recruiting RNs externally provides a “quick fix” to the shortage problem. That is, changing the number of students that are admitted to nursing educational programs will only have an impact once those students have had the chance to complete their programs (either two or four years after their admission to the programs). However, with a lower attrition from the programs (Scenario 2) we also observe a lower initial recruitment of direct care RNs and senior level managers.

**Figure 2.6. Recruitment policies for various scenarios**

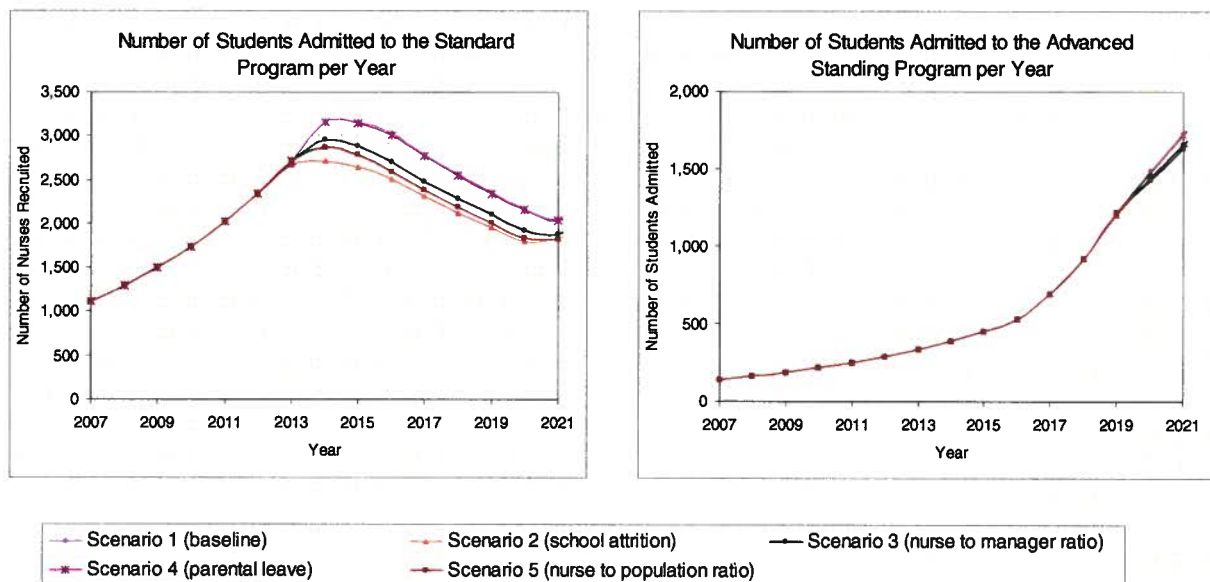


In addition, the model is highly sensitive to the attrition rates of direct care RNs. While Scenario 5 leads to a greater RN population (as seen in Figure 2.5), from Figure 2.6 we note that the higher targets are achieved with lower recruitment needs. This is due to our assumption that an increase in the number of RNs per population will decrease attrition rates of direct care RNs.

We stress that we are not advocating a massive recruitment of RNs in the near future (which has social and ethical implications), but instead are using the model to show that to achieve desired target nurse-population ratios this is necessary. Alternatives could be identified by adding a constraint on the number of RNs recruited early on and relaxing short term RN to population ratios.

Figure 2.7 summarizes the total number of students admitted into BC nursing educational programs. Note that the number of students admitted in the first years of the model is not greatly affected by the parameters analyzed in the scenarios. This is because the capacity of the schools - as represented by physical space, availability of preceptors, clinical placements for students, and faculty - is a limiting factor in our current system (Pringle, Green and Johnson 2006).

**Figure 2.7. Admission policies for various scenarios**




Although the number of students admitted into both types of educational programs must first increase - given that the current programs are not sufficient to meet the assumed demand for RNs, and that a limit on the yearly increase of both types of programs is imposed - after a certain point, the number of students to be admitted into standard programs starts to decrease (up to a certain point) while the number of students to be admitted into the advanced standing program keeps increasing. This occurs despite our assumption that students admitted to the advanced standing program are older than those students admitted to standard programs, and that a fixed cost equivalent to two years of education in a nursing program (in addition to their nursing education

costs) are accrued by every student admitted to the advanced standing program. An explanation of this result might be the lower attrition rates of third year students in relation to first year students or the end effects of a 20 year planning horizon.

Figure 2.8 displays promotion patterns under each scenario. Although the number of RNs promoted from entry-level to senior-level management appears stable, that is not the case for the promotion of direct care RNs to entry-level management. An explanation for the behaviour of the model might be that no limits were imposed on the minimum number of RNs to be promoted. Therefore, because of the insufficient current supply of direct care RNs in the province, the model suggests keeping as many direct care RNs as possible and recruiting to fill entry-level management positions in the first few years. The effect that this has on RN attrition is not known – it may be demoralizing for those who seek career opportunities and must therefore be included in the model.

**Figure 2.8. Promotion policies for various scenarios**



## 2.4 Discussion


The workforce plan might have been affected by the fact that the model was only solved over a finite horizon. For a general discussion of horizon effects in other contexts, see (McClain and Thomas 1977; Fisher, Ramdas and Zheng 2001). While bounds were added to constrain the rate of change of admissions, our solution nonetheless stipulated a decrease in admissions in the last



periods (see Figure 2.9). The size of the effect depended on both the length of the horizon as well as the tightness of the bounds in the last periods.

The analysis above shows that one potential solution to the long term RN shortage is to increase the number of students admitted to an advanced standing educational program while increasing the number of students admitted to standard educational programs at first, but only to a certain point. Although we assumed that students admitted to the advanced standing program are older than those students admitted to standard programs, we also assumed that once they graduate, their attrition from the profession will only depend on their age and not on their type of education. Empirical evidence should be collected to verify this assumption. Furthermore, a large initial recruitment of RNs is required under all scenarios. A possible expansion to the model would consequently be to include a limit on the minimum number of direct care RNs to be promoted or to increase attrition rates of RNs if no such limits are available.

**Figure 2.9. Annual student populations by year of schooling obtained from the optimal solution to the baseline model**



In the first few years, the required number of direct care RNs promoted to entry-level management is small given the expected shortage of direct care RNs. Research has indicated that a sufficient supply of managers is necessary and beneficial for recruitment and retention of direct care RNs (Kramer and Schmalenberg 2002). We could not find research evidence that would point to the “optimal” direct care RN to manager ratio. We used national average ratios of direct care RNs to managers. Further research in this area might be useful to assist decision makers in developing succession plans for managers and senior-level administrators.

Attrition rates vary across nursing education programs (Pringle and Green 2005). In the model, we applied a conservative attrition rate for the nursing educational programs. The model indicated that attrition plays a significant role in the supply of new graduate RNs. An examination of the predictors of successful completion (Patric 2001; Wharrad, Chapple and Price 2003) along with the possible reasons for attrition and strategies to reduce it (Pringle and Green 2005; Scott 2004) may have a significant impact on the entire RN workforce population.

Among the challenges we experienced in applying the model were the limited availability and accessibility of data about the current workforce and the lack of a comprehensive and coordinated data repository or single database on RNs. In 2002, Romanow highlighted that “we cannot expect to keep improving the healthcare system if we do not have the necessary information to measure and track results” (Romanow 2002:xxix). Although data about RNs and physicians are relatively accessible when compared with information about the other health professions, there continue to be significant limitations in the data collected and their availability. These difficulties have been recognized in many reports of the last five years where recommendations or observations regarding the need for improved data collection on HHR are noted (Romanow 2002; Dault et al. 2004; Kephart et al. 2004; Task Force Two 2005; Kephart et al. 2006). We made various assumptions about the model parameters such as direct care RN-to-manager ratios, the minimum number of RNs required in the province per year, and so forth. Future research would help to bring a better empirical basis for the assumptions.

Although the model provides directions, other changes must occur in the system so that model recommendations could be implemented. For example, constraints that limit the capacity of nursing educational programs to expand, such as the shortage of nursing faculty, need to be addressed.

This model could also be used to analyze the impact of policy initiatives such as changes in the composition of multidisciplinary teams and the expanding scope of practice of RNs and licensed practical nurses. Assuming that workforce skill mix will have an impact on the total demand for RNs, this could be accommodated by changing the demand parameters over the years (which are an input to our model) in a similar way as is discussed in Scenario 5. Another possibility would be to incorporate other categories (such as physicians or other nursing groups) in our model. This would require access to the necessary data to populate such a model.



## 2.5 Conclusions and Remarks

We have developed a high level workforce planning model on a provincial level. This model determines the total number of students to admit to programs, the total number of registered nurses to train for management roles, and the total number of nurses and managers resources to recruit from outside the region to meet service needs. The model has been applied to model the registered nurses of British Columbia. We have incorporated factors such as a learning curve for recent graduates and parental leaves (in the calculation of full time equivalence (FTEs)), promotion rules, the existence of two types of educational programs, and the importance of age in the attrition of the workforce.

This model has been formulated as a deterministic linear problem. Given the main goal of our model - to understand how changes in assumptions will impact the optimal number of nurses to train, promote to managerial levels and recruit over the planning horizon - our linear programming approach meets the needs of the decision maker. It provides a framework that can be easily used to generate and compare various scenarios, avoiding the computational burden associated with a stochastic model.

Moreover, additional level of detail may be gained by formulating this problem as a stochastic model. For education, promotion and recruitment policies to be implemented, the effect of variability on the optimal solutions would need to be considered. Possible random variables to be incorporated in the model include: retirement rates, maternity leaves, student attrition, direct care RN-to-manager ratios, and demand for direct care RN. The model would then be formulated in multiple stages: a stage in which these random variables (with assumed joint probability distribution) are observed, and a stage in which our decision variables are optimized given the realization of the random variables discussed. This stochastic program will have much larger dimension than the original deterministic formulation, and further consideration must be given to address the numerical challenges that might be encountered in solving the model (Martel and Price 1981).

We have assumed that the service needs are known. A possible research direction is to combine empirical knowledge of a population's needs and shift scheduling to determine the minimum number of health professionals needed each period. Other extensions to the model include studying the effect of aggregation into age groups, including training of faculty, evaluating whether

a steady state age distribution can be achieved, expanding the model to include other health professionals and directly incorporating feedback loops into the model.

A problem often cited in the literature concerns how changes in the input mix (such as physician-direct care nurse ratios) might impact policy. If our model was expanded to incorporate other health professions, needs for one profession could then be expressed as a function of the needs for the other profession in a similar manner as is done in constraints (2.24) and (2.25). This would also allow us to investigate changes in the scope of practice of each profession by varying the demand parameters, an especially timely subject of discussion (O'Brien-Pallas et al. 2001).

Given the existing concerns regarding the future of the nursing workforce, our model and planning tool provides a systematic way of corroborating that a change in current policy is needed, as well as offering guidance into what policies should be followed to obtain a sustainable workforce.

### CHAPTER 3: A PATIENT SPECIFIC MODEL FOR DECIDING WHEN TO START RADIATION THERAPY FOR PROSTATE CANCER PATIENTS BASED ON THEIR PSA DYNAMICS

Intermediate and high risk prostate cancer patients are often treated with combined hormone therapy and radiation therapy (RT), with some or most of the hormone treatment given before RT. However, the optimal duration of hormone treatment before RT is not known. Finding the best time to initiate RT based on the patient's maximum response to hormone treatment is a key question faced by clinicians (Heymann et al. 2007).

Throughout their hormone treatment, patients are monitored at a few discrete points in time, every one to two months in our study, resulting in a short data series for each patient. By the time sufficient patient-specific data are available, the best time to initiate radiation treatment might already have passed. Similarly, by using an arbitrary duration of hormone treatment before RT, the best time to start RT might not be reached. We present a novel approach to model the disease progression of individual prostate cancer patients on hormone therapy by combining priors based on population characteristics and patient-specific data that are gathered sequentially, so as to make better decisions regarding the best time to initiate RT.

According to the Canadian Cancer Society's Steering Committee (2009), prostate cancer has the highest incidence and is the second leading cause of cancer death in men. In 2009, an estimated 25,500 Canadian men will be diagnosed with prostate cancer and 4,400 will die of the disease. Prostate cancer is characterized by the "uncontrolled growth and spread of abnormal cells" of the prostate (American Cancer Society 2008). Its growth is driven by androgens, such as testosterone, and it is linked to the characteristics of the cells and the number of cells that have the potential to proliferate. Though the disease is initially confined to the prostate, it might spread beyond the prostatic capsule to the tissue surrounding the prostate, the bladder, the pelvic lymph nodes and other parts of the body.

Prostate cancer is typically treated by removing or killing the malignant cells through surgery, radiation, chemotherapy or hormonal therapy (BC Cancer Agency 2009). However, treatments are often combined. For instance, high-intermediate and high risk localized prostate cancer patients, who are scheduled to receive RT, are often treated with hormone therapy beforehand (neoadjuvant hormone therapy). The main goal of the hormone therapy is to starve the tumor cells

of androgens and “induce prostate cancer cell death and tumor regression” (Gleave et al. 2000). By reducing the burden of cancer cells (local, regional or systemic), neoadjuvant hormone therapy might improve the overall effectiveness of the combined treatment (Horwitz and Hanks 2000).

Throughout their hormone treatment, it is common clinical practice to periodically measure patients’ blood levels of prostate specific antigen (PSA). PSA is a protein produced almost exclusively by cells of the prostate. Reference levels for healthy men vary with age, but are generally less than 4 ng/ml in the absence of malignancy. Its levels are known to increase as the volume of the prostate cancer increases (Makarov et al. 2007, Partin and Carter 2000, Sandblom et al. 2002). Since PSA is produced by malignant cells that are growing and dividing, generally the more cancer cells, the more PSA that is produced.

When hormone therapy is given to a patient, we expect to first observe a drastic decrease in the PSA levels in the blood. Later on, the therapy induces a further, but less pronounced, decrease in the PSA levels. This decrease in the PSA levels has been linked to the induction of programmed cell death. In most patients, PSA levels will reach a minimum and then begin to rise. A rise (or insufficient drop) in the PSA values might indicate that the treatment is not being effective or that some of the cells have become resistant to the hormone treatment (Gleave et al. 2000). Such rise in the PSA values during the neoadjuvant hormone treatment has been linked to shorter time to PSA relapse, poorer cause-specific survival, and poorer overall survival (Niblock et al 2006).

Gleave et al. (2000) hypothesized that “maximal tumor regression probably occurs when PSA reaches its nadir level”. We assume that it is ideal to start RT when the PSA reaches its minimum level. However, since cancer progresses and regresses at different rates in patients, it is difficult to predict when such a level will be achieved. Therefore, one standard duration of neoadjuvant androgen ablation (be it 3, 6 or 8 months) is probably not best for everyone (Heymann et al. 2007).

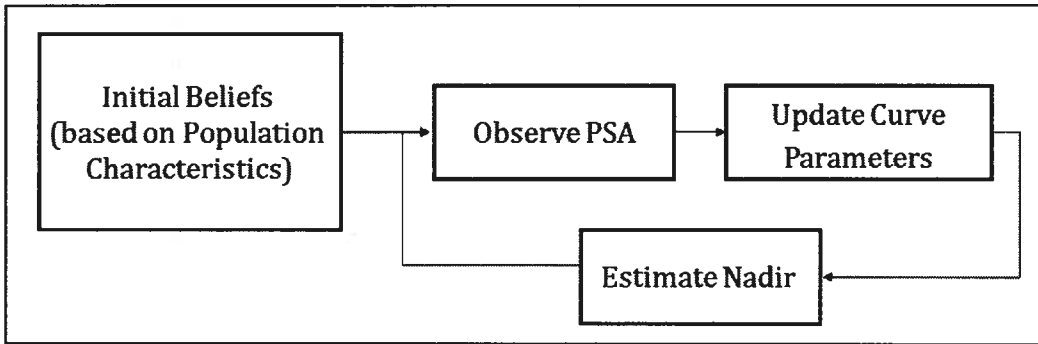
The problem of choosing when to change treatments based on the patient’s specific disease progression is not unique to this context. D’Amato et al. (2000) modeled the decision of when to switch drug therapies based on the virus level of HIV-infected patients. Their goal was to “delay the time until patients are resistant to all existing drug regimens”. In their model, they estimated the probability of switching before the virus reached its minimum level and the mean delay in detection of viral rebound. They tested two policies using simulation: the *viral load policy*, in which therapies were switched if the current viral load was higher than the previous minimum viral load

by a given threshold, and the *proactive policy*, in which patients switched treatments if the viral load's threshold was reached or a predetermined switching time elapsed.

Our approach differs from the one described above in that we consider two sources of randomness: the variability associated with each PSA reading, and the variability of PSA progression. The PSA dynamics are modeled as a log quadratic function of the time elapsed since hormone therapy is initiated. Our decision to start RT is based on the probability of choosing the time when the PSA reaches its minimum level. Key to our model is the incorporation of parameter updating as each new PSA reading becomes available.

The main steps involved in our modeling approach are summarized in Figure 3.1. We start by modeling the initial beliefs of response patterns based on population characteristics. As new information becomes sequentially available, we use the Kalman filter to update the estimates of the curve parameters. We determine the distribution of the time of the nadir given our updated curve parameters. As new PSA values are observed, we use the Kalman filter recursively to update the curve parameters and the distribution of the nadir.

**Figure 3.1. Model overview.**



This chapter is organized as follows. The PSA dynamics model is described in Section 3.1.1. Section 3.1.2 discusses how the prior distribution of the curve parameters of the PSA dynamics model is estimated. The parameters are used to estimate the time of the PSA nadir, which is linked to the optimum RT time, in Section 3.1.3. Section 3.1.4 describes the state space model used to sequentially update the distribution of the parameters as new PSA readings become available. More specificity is obtained by clustering patients in Section 3.1.5. We illustrate our modeling approach in Section 3.2 by comparing clinically implementable policies on a cohort of intermediate risk prostate cancer patients who were enrolled in a prospective randomized trial (Morris et al. 2009). We conclude with final remarks in Section 3.3.

### 3.1 Methods

Table 3.1 summarizes the notation used in our model.

Table 3.1. Model notation.

Notation	Description
$i$	Patient index
$j$	Cluster index
$t$	Time index
$tnadir_i$	Estimated time of nadir from the time at which hormone therapy started
$minPSA_i$	Minimum PSA value observed
$N$	Number of patients in the population
$M$	Number of clusters
$Y_{i,t}$	Observation variable (log of the prostate specific antigen level)
$\theta_{i,j,t} = (\alpha_{i,j,t}, \beta_{i,j,t}, \gamma_{i,j,t})^T$	True regression parameters
$\bar{\theta}_j = (\bar{\alpha}_j, \bar{\beta}_j, \bar{\gamma}_j)^T$	Prior estimates of the regression parameters
$\bar{\theta}_j^w = (\bar{\alpha}_j^w, \bar{\beta}_j^w, \bar{\gamma}_j^w)^T$	
$\hat{\theta}_i = (\hat{\alpha}_i, \hat{\beta}_i, \hat{\gamma}_i)^T$	Estimated regression coefficients given all PSA readings from the start of the hormone therapy up to the start of RT
$\hat{\theta}_{i,j,t t-1}$	Kalman filter estimate of $\theta_{i,j,t}$ given all PSA readings up to time $t-1$
$\varepsilon_{i,j,t}$	Observation error
$V_{i,j,t}$	Variance of observation error
$R_{i,j,t}$	Between patient variance of the estimates
$R_{i,j,t t-1}$	Between patient variance of the estimates given all PSA readings up to time $t-1$
$\omega_{i,j,t}$	Parameter errors
$W_{i,t}$	Covariance of parameter errors
$P(j)_{i,t}$	Probability of an observation belonging to the $j^{th}$ cluster
$\rho_j$	Mean vector and the covariance matrix for cluster $j$

#### 3.1.1 Modeling PSA Dynamics

The PSA progression of patient  $i$  as a function of the time  $t$  elapsed since the patient started his hormone treatment (in days) can be modeled by a log quadratic curve:

$$Y_{i,t} = \ln(PSA_{i,t}) = \alpha_i + \beta_i t + \gamma_i t^2 + \varepsilon_i \quad \varepsilon_i \sim N(0, V_i) \quad (3.1)$$

Equation (3.1) can be justified in two ways: theoretically and empirically. Assume that the size of the tumor  $N_i(t)$ , represented by the number of tumor cells, is proportional to the levels of PSA in the blood  $PSA_{i,t}$ . To model the dynamics of the PSA values as a function of time, it is first necessary to understand the impact that the hormone treatment has on the tumor size. Based on Goldie and Coldman's (1979) hypothesis, the size of the tumour may be expressed as the sum of the resistant and the non resistant cell population. For each patient  $i$ , let  $X_i(t)$  be the number of tumor cells at time  $t$  that depend on androgens to grow, and let  $Z_i(t)$  be the number of cells that can grow and divide in the absence of androgens (e.g., androgen independent cells). The total size of the tumor  $N_i(t)$  can therefore be expressed as the sum of  $X_i(t)$  and  $Z_i(t)$ .

At time  $t$ , there are two competing factors that affect the size of the tumor of patient  $i$ : cell division  $g_i$  and cell death  $a_i$ . Under sufficient androgen levels, we assume that the growth rate of the tumor is a linear function of its size, that is:

$$dN_i(t)/dt = (g_i - a_i) N_i(t) \quad (3.2)$$

By solving (3.2),  $N_i(t)$  can be written as  $C_i \exp((g_i - a_i)t)$ , where  $C_i$  is a constant that depends on the patient specific characteristics.

When hormone therapy is given, its main impact is to reduce the levels of androgens or their effectiveness. Only androgen independent cells are able to divide, while all cells die due to apoptosis, a natural process of self-destruction in certain cells. We assume that the ability of androgen independent cells to divide increases linearly over time. The growth rate is given by

$$dZ_i(t)/dt = g_i t Z_i(t) - a_i Z_i(t) \quad (3.3)$$

and

$$dX_i(t)/dt = -a_i X_i(t) \quad (3.4)$$

If we let  $m_i$  be the fraction of androgen independent cells, and assume that the number of androgen independent cells is proportional to the total number of cells,  $Z_i(t)$  can be rewritten as  $Z_i(t) = m_i N_i(t)$ . From equations (3.3) and (3.4), the net growth rate can be rewritten as:


$$dN_i(t)/dt = g_i m_i t N_i(t) - a_i N_i(t) \quad (3.5)$$

By solving (3.5), we obtain  $N_i(t) = C_i \exp(1/2g_i m_i t^2 - a_i t)$ . Assuming that  $N_i(t) \propto PSA_{i,t}$ , and under the assumption of normality of the errors of the PSA readings, the PSA dynamics can therefore be modeled as:

$$Y_{i,t} = \ln(PSA_{i,t}) = \alpha_i + \beta_i t + \gamma_i t^2 + \varepsilon_i \quad \varepsilon_i \sim N(0, V_i) \quad (3.6)$$

A typical curve appears in Figure 3.2. It provides some empirical justification for this model.  $R^2$  values for the 163 patients on our study range from 0.57 to 0.99 with an average  $R^2$  of 0.9. Further analysis of the goodness of fit of the curve and how  $R^2$  varies by patient type is presented in Chapter 4.

**Figure 3.2. PSA versus time curve fit for a randomly chosen patient.**



The axes values have been omitted because of patient confidentiality. Note that the y-axis is plotted on a logarithmic scale. The curve is obtained by fitting the regression curve  $Y_{i,t} = \alpha_i + \beta_i t + \gamma_i t^2 + \varepsilon_i$ ,  $\varepsilon_i \sim N(0, V_i)$ . For comparison purposes, PSA values observed are included in the graph.

### 3.1.2 Estimating the Prior Distribution of the Curve Parameters

Assume each patient's curve parameters are drawn from a multivariate normal distribution with mean  $\bar{\theta}$  and variance  $R$ . To estimate the distribution of the parameters in the population, we start by fitting equation (3.1) for each patient  $i$  given all PSA readings from the start of the hormone therapy up to the time  $RT$  is given. The estimated regression coefficients are denoted  $\hat{\theta}_i = (\hat{\theta}_{1,i}, \hat{\theta}_{2,i}, \hat{\theta}_{3,i})^T = (\hat{\alpha}_i, \hat{\beta}_i, \hat{\gamma}_i)^T$ .



Our next step is to estimate the priors of the PSA dynamics of new patients that initiate hormone therapy. An estimate of the initial beliefs of the regression coefficients  $\bar{\theta} = (\bar{\theta}_1, \bar{\theta}_2, \bar{\theta}_3)^T = (\bar{\alpha}, \bar{\beta}, \bar{\gamma})^T$  can be obtained based on the mean of the regression coefficients of all patients in the population:

$$\bar{\theta}_k = \sum_{i=1}^n \hat{\theta}_{k,i} / n \text{ for } k = 1, 2, 3 \quad (3.7)$$

Alternatively, regression coefficients can also be weighted by their variance  $Var(\hat{\theta}_i) = (Var(\hat{\theta}_{1,i}), Var(\hat{\theta}_{2,i}), Var(\hat{\theta}_{3,i}))^T = (Var(\hat{\alpha}_i), Var(\hat{\beta}_i), Var(\hat{\gamma}_i))^T$  to obtain a more accurate estimate of the regression coefficients of new patients. Coefficients with smaller variability receive more weight than coefficients with greater variability according to:

$$\bar{\theta}_k^w = \sum_{i=1}^n \frac{\hat{\theta}_{k,i}}{Var(\hat{\theta}_{k,i})} / \sum_{i=1}^n \left( \frac{1}{Var(\hat{\theta}_{k,i})} \right) \text{ for } k = 1, 2, 3 \quad (3.8)$$

We determine the between patient variance of the estimates by calculating the covariance of the regression coefficients of all patients in the population. The variance of the PSA readings  $V_i$  is calculated as the average of the mean square errors of all patients in the population. Other alternatives may be explored using parametric empirical Bayes' inferences to improve our estimation of the parameters (Morris 1983).

### 3.1.3 Modeling the Distribution of the Nadir

By taking the derivative of the Equation (1) and setting it to zero, we are able to estimate  $tnadir_i$  (the time at which patient  $i$  reaches his PSA nadir) by:

$$tnadir_i = -\beta_i / (2\gamma_i) \quad (3.9)$$

For  $tnadir_i$  to be well defined, we must assume that  $\gamma_i \neq 0$ . In addition, we assume that the nadir is reached after hormone therapy is given ( $tnadir_i > 0$ ). As indicated by the BCCA's current protocol (BC Cancer Agency 2009), RT must be started at the end of eight months (at time  $t = 240$ ) unless a nadir has been reached beforehand. Equation (3.9) can therefore be rewritten as:


$$tnadir_i = \begin{cases} \min(-\beta_i/(2\gamma_i), 240) & \gamma_i * \beta_i < 0 \\ 240 & \text{otherwise} \end{cases} \quad (3.10)$$

where  $(\beta_i, \gamma_i) \sim N\left((\hat{\beta}_i, \hat{\gamma}_i); \begin{pmatrix} r_{22i} & r_{23i} \\ r_{32i} & r_{33i} \end{pmatrix}\right)$ .

Note that a faster growth rate is linked to a higher number of tumor cells, which will lead to a higher rate of cells dying. Therefore,  $\beta_i$  and  $\gamma_i$  are negatively correlated.

As can be seen in Figure 3.3, it is important to incorporate the correlation of the parameters in our distribution calculations. The correlation can significantly impact the estimation of when the maximum nadir probability is reached and our decision to start radiation therapy.

**Figure 3.3. Impact of incorporating correlation in the distribution of the time of nadir for a randomly chosen patient.**



The distribution under the assumption of independence and correlation are obtained by simulating 10,000 observations for the given patient assuming that the parameters are normally distributed with mean  $(\hat{\beta}_i, \hat{\gamma}_i)$  and covariance matrix  $\begin{pmatrix} r_{22i} & 0 \\ 0 & r_{33i} \end{pmatrix}$  and  $\begin{pmatrix} r_{22i} & r_{23i} \\ r_{32i} & r_{33i} \end{pmatrix}$  respectively.

The distribution of the ratio of correlated normal random numbers was studied by Fieller (1932). Following the results presented by Hinkley (1969), we show that, as  $P(\gamma_i > 0) \rightarrow 1$ , the cumulative distribution of the time of the nadir converges to:

$$F(t_{nadir}) = \frac{1}{\sqrt{2\pi}} \left( \frac{(\hat{\beta}_i + 2\hat{\gamma}_i t)}{\left( 2\sqrt{r_{33i}}\sqrt{r_{22i}}\sqrt{\frac{t^2}{r_{22i}} + \frac{r_{32i}t}{r_{22i}r_{33i}} + \frac{1}{4r_{33i}}} \right)} \int_{-\infty}^{\cdot} \exp\{-y^2/2\} dy \right) \quad (3.11)$$

Let the time of nadir of patient  $i$  be given by  $-\beta_i/(2\gamma_i)$ , where  $(\beta_i, \gamma_i) \sim N\left((\hat{\beta}_i, \hat{\gamma}_i); \begin{pmatrix} r_{22i} & r_{23i} \\ r_{32i} & r_{33i} \end{pmatrix}\right)$ .

Note that:  $E(-\beta_i) = -\hat{\beta}_i$ ,  $E(2\gamma_i) = 2\hat{\gamma}_i$ ,  $Var(-\beta_i) = r_{22i}$ ,  $Var(2\gamma_i) = 4r_{33i}$ ,  $Cov(-\beta_i, 2\gamma_i) = -2r_{32i}$  and  $Corr(-\beta_i, 2\gamma_i) = -r_{32i}/(\sqrt{r_{33i}}\sqrt{r_{22i}})$ . Let

$$g(t) = (2r_{33i}t + r_{32i}) / \left( 2r_{33i}\sqrt{r_{22i}}\sqrt{\frac{t^2}{r_{22i}} + \frac{r_{32i}t}{r_{22i}r_{33i}} + \frac{1}{4r_{33i}}} \right) \quad (3.12)$$

and

$$h(t) = (-\hat{\beta}_i - 2\hat{\gamma}_i t) / \left( 2\sqrt{r_{33i}}\sqrt{r_{22i}}\sqrt{\frac{t^2}{r_{22i}} + \frac{r_{32i}t}{r_{22i}r_{33i}} + \frac{1}{4r_{33i}}} \right) \quad (3.13)$$

Based on the results by Fieller (1932), the cumulative distribution of the time of nadir is given by:

$$\frac{1}{2\pi\sqrt{1-g(t)^2}} \left( \int_{h(t)}^{\infty} \int_{\frac{-\hat{\gamma}_i}{\sqrt{r_{33i}}}}^{\infty} \exp\left\{\frac{-x^2 + 2g(t)xy - y^2}{2(1-g(t)^2)}\right\} dx dy + \int_{-h(t)}^{\infty} \int_{\frac{\hat{\gamma}_i}{\sqrt{r_{33i}}}}^{\infty} \exp\left\{\frac{-x^2 + 2g(t)xy - y^2}{2(1-g(t)^2)}\right\} dx dy \right) \quad (3.14)$$

Letting  $P(\gamma_i > 0) \rightarrow 1$  (or  $\hat{\gamma}_i/\sqrt{r_{33i}} \rightarrow \infty$ ), the second term in equation (3.14) approaches 0. Given

that  $\int_{-\infty}^{\infty} \exp\{-ax^2 - 2bx\} dx = \sqrt{\pi/a} \exp\{b^2/a\}$ , the first term in equation (3.14) approaches:

$$\frac{\sqrt{2\pi(1-g(t)^2)}}{2\pi\sqrt{1-g(t)^2}} \left( \int_{h(t)}^{\infty} \exp\{-y^2/(2(1-g(t)^2))\} \exp\{(g(t)^2 y^2)/(2(1-g(t)^2))\} dy \right).$$

This simplifies to:

$$\frac{1}{\sqrt{2\pi}} \left( \int_{h(t)}^{\infty} \exp\{-y^2/2\} dy \right) \quad (3.15)$$

Which is equivalent to:

$$\frac{1}{\sqrt{2\pi}} \left( \int_{-\infty}^{-h(t)} \exp\{-y^2/2\} dy \right) \quad (3.16)$$

Note that, following a similar argument, it is possible to prove that, as  $P(\gamma_i < 0) \rightarrow 1$ , the distribution of the time of nadir converges to:

$$\frac{1}{\sqrt{2\pi}} \left( \int_{-\infty}^{h(t)} \exp\{-y^2/2\} dy \right) \quad (3.17)$$

Results presented by Hinkley (1969) can also be used to obtain bounds on this approximation. Letting  $F^*(t_{nadir})$  be the true cumulative distribution of the time of the nadir and  $F(t_{nadir})$  be the approximation provided by (3.11):

$$|F(t_{nadir}) - F^*(t_{nadir})| \leq 1 - \Phi \left( r_{32_i} / \sqrt{r_{33_i}^* r_{22_i}} \right) \quad (3.18)$$

Where  $\Phi$  is the density function of a standard normal random variable.

We compared this approximation to the results obtained by simulating this distribution. For each patient, we simulated 10,000 observations with mean  $(\hat{\beta}_i, \hat{\gamma}_i)$  and covariance matrix  $\begin{pmatrix} r_{22_i} & r_{23_i} \\ r_{32_i} & r_{33_i} \end{pmatrix}$ . Using equation (3.10), we estimated the time of nadir for each observation. By discretizing time, we estimated the proportion of times (out of 10,000) that the simulated nadir fell into each time category. As can be seen in Section 3.2.2, the two methods of estimating the period of maximum probability gave similar results, suggesting that this approximation is appropriate.

### 3.1.4 Updating the Parameter Estimation Using a State Space Model

Given the short data series for each patient, we include prior information to augment the data and gain knowledge of what to expect of the population. We aim to combine our prior knowledge of the parameters with the new observations that become sequentially available. For this purpose, we use the Kalman filter to update the estimates of the curve parameters as new information is obtained. An alternative to ordinary least squares regression, the Kalman filter is a *recursive* procedure that computes the optimal estimator of the state vector at each time period based on a series of noisy measures (Harvey 1991). While results obtained from the Kalman Filter model may be derived from Multivariate Normal results, considering this model in the state space framework provides the additional value of setting the problem for dynamic programming calculations to be performed in the future.

The Kalman filter has been used in a wide variety of patient specific models. Some examples include: tracking urinary bladder filling (Kristiansen et al. 2005), monitoring glucose levels in the blood (Parker et al. 1999, Knobbe and Buckingham 2005) and tracking the position of a tumor in patients receiving radiation therapy (Rehbinder et al. 2004). In our model, the noisy measures correspond to the PSA readings and the state vector corresponds to the PSA dynamics.

We start by putting the model in a state space form. For that purpose, we determine the observable variables and how they are related to the (non-observable) parameters of interest. Let  $Y_{i,t}$ ,  $\varepsilon_{i,t}$  and  $V_{i,t}$  be scalars representing the observable variable ( $\ln(PSA)$ ), the observation error and the variance of the observation error of patient  $i$  at time  $t$ . Let  $F_{i,t}$  be the  $3 \times 1$  vector  $(1, t, t^2)^T$  where  $t$  is the time elapsed from the start date of the hormone treatment for patient  $i$ . Let  $\theta_{i,t}$  be the  $3 \times 1$  vector  $(\alpha_{i,t}, \beta_{i,t}, \gamma_{i,t})^T$  containing the curve parameters for patient  $i$  at time  $t$ . The observable variable  $Y_{i,t}$  is related to the curve parameters  $\theta_{i,t}$  by the measurement equation:

$$Y_{i,t} = F_{i,t}^T \theta_{i,t} + \varepsilon_{i,t} \quad \varepsilon_{i,t} \sim N(0, V_{i,t}) \quad (3.19)$$

We assume that the errors are temporally and mutually independent and that the disturbances and the initial state vector are normally distributed. In addition, we assume that the underlying model does not change over time. That is,  $\theta_{i,t} = \theta_{i,s}$  for all  $t, s$ . We can therefore write:

$$\theta_{i,t} = \theta_{i,t-1} + \omega_{i,t} \quad \omega_{i,t} \sim N(0, W_{i,t}) \quad (3.20)$$

where  $\omega_{i,t}$  is the 3x1 vector of the parameters' errors with mean zero and covariance matrix  $W_{i,t}$ . Given our assumption that the true model does not change over time, the matrix  $W_{i,t}$  is assumed to have all components equal to 0 at  $t > 0$ . If measurement error is incorporated in the initial estimation of the parameters, this is done by letting  $W_{i,0}$  be the covariance matrix of the initial parameters.

Now we apply the Kalman filter to obtain estimates of the parameters of the above model. Let  $\hat{\theta}_{i,t|t-1}$  be the estimate of  $\theta_{i,t}$  of patient  $i$  at time  $t$  given all PSA readings for that patient up to time  $t-1$ . When a new PSA reading,  $Y_{i,t}$ , becomes available,  $\hat{\theta}_{i,t|t-1}$  is updated as follows:

$$\hat{\theta}_{i,t|t} = \hat{\theta}_{i,t|t-1} + R_{i,t|t-1} F_{i,t} Q_{i,t}^{-1} [Y_{i,t} - F_{i,t}^T \hat{\theta}_{i,t|t-1}] \quad (3.21)$$

where  $R_{i,t|t}$  is the 3x3 covariance matrix of the estimation error.  $R_{i,t|t}$  is updated as follows:

$$R_{i,t|t} = R_{i,t|t-1} - R_{i,t|t-1} F_{i,t} Q_{i,t}^{-1} F_{i,t}^T R_{i,t|t-1} \quad (3.22)$$

$$Q_{i,t} = F_{i,t}^T R_{i,t|t-1} F_{i,t} + V_{i,t} \quad (3.23)$$

From equation (3.20),

$$\hat{\theta}_{i,t|t-1} = \hat{\theta}_{i,t-1|t-1} \quad (3.24)$$

and

$$R_{i,t|t-1} = R_{i,t-1|t-1} + W_{i,t} \quad (3.25)$$


Furthermore, observe that  $Y_{i,t}$  enters only in equation (3.21). Since the covariance of the parameter estimates does not depend on the data obtained in each period, it can be computed off-line based on the prior information.

### 3.1.5 Clustering

Finding the prior distribution for the parameters can be challenging. As illustrated by Zhu and Rohwer (1996), prior distribution assumptions are key to any Bayesian framework. The clinicians in our team believe that certain patients may be faster or slower responders to hormone therapy. By clustering, as opposed to using one grouping of the entire population of data, we can separate

out this variability and tailor our model to these different patient subgroups. Figure 3.4 shows how clustering can be used to improve our model estimates.


**Figure 3.4. Main steps in clustering the prior distribution.**



As in Section 3.1.2, we first fit regression curve (3.1) for each patient  $i$ . The distribution of  $tnadir_i$  as well as the minimum PSA value observed ( $minPSA_i$ ) for each patient is calculated. Other possible clustering variables include the regression parameters. However, as is discussed later, clustering based on  $tnadir_i$  and  $minPSA_i$  gives results that are easiest to interpret.

As can be observed in Figure 3.5 and Figure 3.9, the population is not homogeneous: there appears to be subgroups within the population.

**Figure 3.5. Distribution of the estimated time to reach the nadir and the minimum PSA value observed.**



We therefore partition patients into subgroups based on the  $tnadir_i$  and  $minPSA_i$ . For this purpose, we use the model based strategy for clustering proposed by Fraley and Raftery (1998). This method combines agglomerate hierarchical clustering with the EM algorithm to maximize the likelihood that observation  $x_i$  belongs to group  $j$ :

$$L(\rho_1, \dots, \rho_M; P(1), \dots, P(M)) = \prod_{i=1}^n \sum_{j=1}^M P(j) f_j(x_i | \rho_j) \quad (3.26)$$

where  $M$  represents the number of clusters,  $P(j)$  is the probability of an observation belonging to the  $j^{th}$  cluster,  $\rho_j$  are the mean vector and the covariance matrix for cluster  $j$ , and  $f_j(x_i | \rho_j)$  is the density of each observation (assumed to be multivariate normal).

In our model,  $\rho_j$  corresponds to the mean and covariance matrix of  $tnadir_i$  and  $minPSA_i$  for all patients in cluster  $j$ . We begin by letting each observation be in a cluster by itself. Sequentially, observations are merged greedily based on (3.26) until  $M$  clusters are formed. Using this partition to initialize the EM algorithm, we iterate between an M step in which maximum-likelihood parameter estimates given the clustering partition are computed and an E step in which a clustering partition is obtained based on the maximum likelihood parameter estimates until the relative difference between successive parameter estimates is below a threshold. Each observation is taken to be part of the group that has the maximum conditional probability.

Once observations are classified, we study the distribution of the regression parameters within each cluster  $j$ . Equations (3.7) and (3.8) can be rewritten as:

$$\bar{\theta}_{k,j} = \sum_{i=1}^n c_{i,j} \hat{\theta}_{k,i,j} / \sum_{i=1}^n c_{i,j} \quad \text{for } k = 1, 2, 3 \quad (3.27)$$

and

$$\bar{\theta}_{k,j}^w = \sum_{i=1}^n \frac{1}{Var(\hat{\theta}_{k,i,j})} (c_{i,j} \hat{\theta}_{k,i,j}) / \sum_{i=1}^n \left( \frac{1}{Var(\hat{\theta}_{k,i,j})} c_{i,j} \right) \quad \text{for } k = 1, 2, 3 \quad (3.28)$$

where  $c_{i,j}$  equals 1 if patient  $i$  belongs to cluster  $j$ , 0 otherwise.

We determine the between patient variance of the estimates by calculating the covariance of the regression coefficients of all patients in the cluster. The variance of the observations is assumed to be the average of the mean square errors of all patients in the cluster.

Our next step is to estimate the probability that a new patient belongs to a cluster given the patient specific characteristics such as his initial PSA level, co-morbidities and staging (AJCC 2002). We opted for a logit model (Agresti 1996) given its simplicity to interpret. The initial beliefs are based



on the parameter estimates of the cluster for which the patient has the highest probability of belonging.


We can proceed in two ways. One way is to assign patients initially to a cluster and only update the parameters for that cluster as new information becomes available. Another possibility would be to update the probability that patient  $i$  belongs to cluster  $j$  at time  $t$  ( $P(j)_{i,t}$ ) as new PSA readings become available. In each period  $t$ , the patient would either be assigned the cluster with the highest probability or a weighted average of the nadir probability for all clusters would be taken. While this entails keeping track of the curve parameters for each cluster at all time periods, this might, on the other hand, improve our initial cluster probability estimates for patient  $i$ .  $P(j)_{i,t}$  is updated as follows:

$$P(j)_{i,t} = \frac{P(j)_{i,t-1} * f(PSA_{i,t} | (F_{i,t}^T \hat{\theta}_{i,j,t|t-1}; Q_{i,j,t-1}))}{\sum_{j=1}^3 (P(j)_{i,t-1} * f(PSA_{i,t} | (F_{i,t}^T \hat{\theta}_{i,j,t|t-1}; Q_{i,j,t-1})))} \quad (3.29)$$

where  $f(y | (\mu; \sigma^2))$  is the univariate normal density function with mean  $\mu$  and variance  $\sigma^2$ .

Equation (3.29) states that the probability that a patient is in cluster  $j$  given all the information available up to time  $t$  is proportional to the prior probability that the patient belongs to such group times the probability that a PSA value  $PSA_{i,t}$  is observed at time  $t$ . Figure 3.6 summarizes the parameter updating process.

**Figure 3.6. Steps involved in the curve parameter updating**



### 3.2 Implementation and Data Analysis

We implemented our modeling approach on data from 163 intermediate risk patients who were part of a prospective randomized trial (Morris et al. 2009). The purpose of the clinical trial was to

study the impact on biochemical recurrence of RT dose escalation using implanted radioactive iodine sources, compared with dose escalated external beam RT. Prior to dose escalation, all of the patients received 8 months of luteinizing hormone-releasing therapy, with at least one month of nonsteroidal antiandrogen hormones, followed by pelvic external beam radiation therapy. Usually patients had PSA readings every two months before radiotherapy, as was specified in the study protocol. Over 70% of the patients had at least five PSA readings before radiotherapy.

For each patient, data available included: hormone and radiation therapy start date as well as the dates and values of the PSA readings taken during hormone treatment. Additional information available for each patient included: the type of drugs given during hormone treatment, whether the patient switched the type of hormone treatment, the dates and values of testosterone readings, whether the patient was diabetic, had vascular disease or had other bilateral diseases, the initial T stage, Gleason grade of cancer score and the percent of biopsies that contained cancer.

We now describe the analysis performed when incorporating clustering to improve our model estimates. In an implementation that leaves out clustering assumptions, all clustering steps would be omitted, and all members of the population would be taken to be part of a single cluster. Policy results obtained under both sets of assumptions are discussed in Section 3.2.2.


### 3.2.1 Data Analysis

Looking retrospectively, we fit a regression curve for each patient in our population based on all the PSA observations obtained for that patient from the time at which hormone therapy started (time 0). To test goodness of fit, we compared the predicted time of the nadir ( $tnadir_i$ ) obtained by fitting the regression curve (1) for each patient to the time at which the minimum PSA was observed for that patient (e.g., the time of  $minPSA_i$ ). This was done by performing a paired data t-test. We tested the null hypothesis that the times were equal versus the alternative hypothesis that the times differed. There was not sufficient evidence to reject our null hypothesis ( $t = 1.13$ ,  $P > .25$ ). Note that PSA readings were only taken at discrete points in time (every one to two months). The time at which the readings were taken might impact the time at which the minimum PSA  $minPSA_i$  was observed.

Figure 3.7 shows a scatterplot of the regression parameters obtained for all patients. As can be seen in Figure 3.7, the curve parameters  $\hat{\beta}_i$  and  $\hat{\gamma}_i$  are negatively correlated (the correlation

coefficient between  $\hat{\beta}_i$  and  $\hat{\gamma}_i$  is -0.92). This supports our earlier remark that omitting the parameter correlation in the nadir calculation is not appropriate.

Figure 3.7.  $\hat{\beta}_i$  and  $\hat{\gamma}_i$  from the regression curve  $Y_{i,t} = \alpha_i + \beta_i t + \gamma_i t^2 + v_i$  calculated for each patient in the dataset.




Patients have been clustered based on  $tnadir_i$  and  $minPSA_i$ .

It was not easy to choose which variables to use as the basis for clustering. However, the expected time to reach the nadir  $tnadir_i$  and the minimum PSA value observed prior to the start of the radiation treatment  $minPSA_i$  gave results that were easiest to interpret. Other possible clustering variables considered included the regression parameters  $\hat{\alpha}_i$ ,  $\hat{\beta}_i$  and  $\hat{\gamma}_i$ , the estimated minimum PSA value (estimated in a similar way as  $tnadir_i$ ) and the initial PSA. Note that, as part of the selection criteria for the clinical study we applied our analyses to, the group of patients was relatively homogeneous - primarily intermediate and lower tier high risk prostate cancer patients. In a more heterogeneous group, we would expect other explanatory variables to play a role in our clustering model and we are investigating this further.

The number of clusters was chosen based on the Bayesian Information Criterion values (Fraley and Raftery, 1998) as well as on the ease of interpretation of the results. The Bayesian Information Criterion (BIC) values as a function of the number of clusters  $M$  in the population is presented in Figure 3.8. Based on the best BIC values obtained within each clustering assumption, clustering the

population into either three or four subgroups is suggested. While clustering patients into four groups gives a slightly higher BIC value than clustering into three groups (a BIC of -1639.32 versus -1642.21), we chose three clusters given the ease of interpretation by clinicians. Other possible criteria to be considered when assessing the number of clusters are described by (Celeux and Soromenho, 1996).

**Figure 3.8. Summary of best BIC as a function of the number  $M$  of clusters in the population.**



The best BIC for each  $M$  is obtained by comparing the BIC under various possible parameterizations of the covariance matrix (Fraley and Raftery, 1998).

Table 3.2 compares average cluster probabilities given the most likely cluster membership. While, as expected, the average probabilities are higher for the diagonals, patients are not perfectly separated.

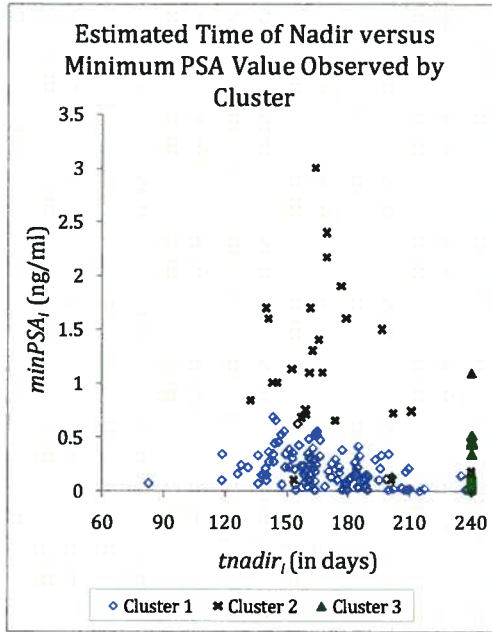
**Table 3.2. Average cluster probabilities given most likely cluster membership.**

Most likely cluster/Possible clusters	Cluster 1	Cluster 2	Cluster 3
Cluster 1	0.49	0.31	0.20
Cluster 2	0.29	0.39	0.25
Cluster 3	0.21	0.27	0.52

The most likely cluster represents the cluster to which each patient is classified. Within all patients in a given cluster, we compute the average probability of belonging to each of the possible clusters. Higher probabilities of the diagonals represent better cluster separation. Note that, since we are dealing only with averages, neither the rows nor the columns in the table necessarily add to one.

Figure 3.9 summarizes the clusters obtained. Note that 72% of the patients were assigned to cluster 1, which has an average estimated time to reach the nadir that is below the current guidelines of 8 months. Patients in the second cluster have an average time to reach the nadir of around 8 months and patients in the third cluster have a longer time to reach the nadir.

**Figure 3.9. Cluster classification of patients based on the estimated time of nadir and minimum PSA value observed.**

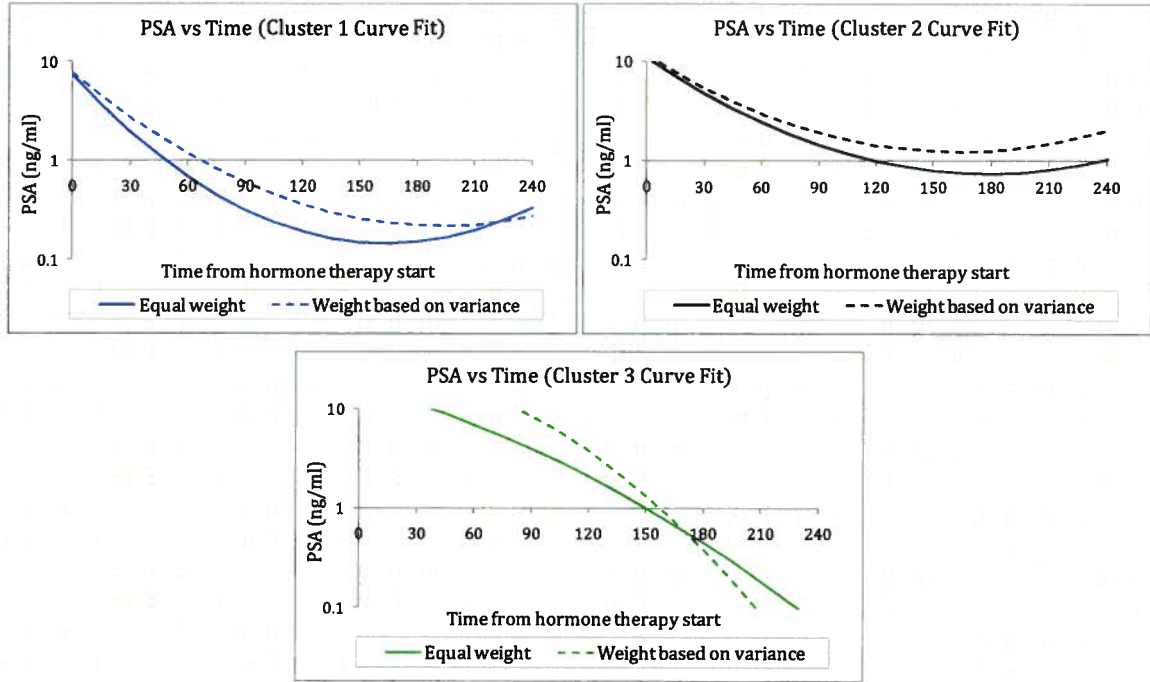


	Cluster 1	Cluster 2	Cluster 3
<i>tnadir</i> (mean, in months)	5.5	8	>> 8
<i>minPSA</i> (mean [variance])	.21[.03]	1.03[.56]	.26[.09]
Proportion of patients in cluster	72%	20%	8%

From Figure 3.7, note that patients in the third cluster are the only patients to present a  $\hat{\gamma}_i < 0$ . This third cluster is comprised mostly of patients whose initial hormone treatment strategy was not effective and had to be given additional hormones during treatment. Giving additional hormones adds cost and morbidity to the treatment. Being able to predict which patients are in the third cluster earlier, and which patients will therefore require additional hormones, is important to clinicians.

Using equations (3.27) and (3.28), we estimated the parameters within each cluster. The results obtained are presented in Figure 3.10.

**Figure 3.10. Regression curves based on prior means obtained by weighting parameters within each cluster.**



Parameters are weighted based on Equation (3.27) (equal weight) or on Equation (3.28) (weight based on variance).

The impact on policy decisions of using either equation to estimate the initial parameters within each cluster is shown in Section 3.2.2.

Finally, we estimated the probability of each patient being in a cluster using the following logit model (Agresti 1996):

$$P(1)_{i,0} = \exp(2.3 - 0.16 \cdot PSA_{i,0}) / [1 + \exp(2.3 - 0.16 \cdot PSA_{i,0}) + \exp(1.16 - 0.06 \cdot PSA_{i,0})] \quad (3.30)$$

$$P(2)_{i,0} = \exp(1.16 - 0.06 \cdot PSA_{i,0}) / [1 + \exp(2.3 - 0.16 \cdot PSA_{i,0}) + \exp(1.16 - 0.06 \cdot PSA_{i,0})] \quad (3.31)$$

$$P(3)_{i,0} = 1 - P(1)_{i,0} - P(2)_{i,0} \quad (3.32)$$

$P(j)_{i,0}$  represents the probability that a patient  $i$  belongs to cluster  $j$  at the beginning of the neoadjuvant hormonal treatment (e.g. at time 0).  $PSA_{i,0}$  is the PSA value for patient  $i$  at time 0.

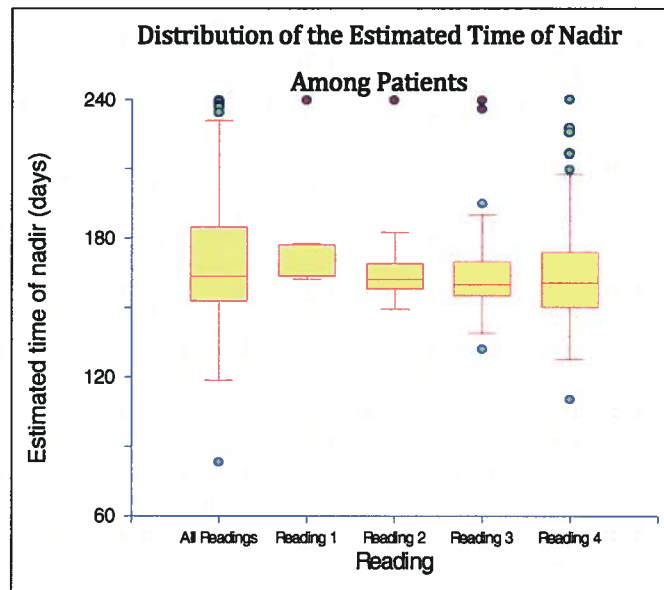
Baseline values considered included: initial PSA value, Gleason score, T stage, percent of biopsies that contained cancer, whether the patient was diabetic and whether he had any vascular disease.

Nevertheless, when the initial PSA value was used, none of the other variables was a significant predictor. As additional PSA values became available for a given patient, we used the Kalman filter to update the regression parameter estimates. Based on Equation (3.10), we then estimated the time of nadir.

Figure 3.11 summarizes how the distribution of the estimated time of nadir between patients changes given the number of times the parameters for each patient are updated (Reading 1, Reading 2, Reading 3 and Reading 4). We compare it with the distribution of the estimated time of nadir ( $tnadir_i$ ) that does not incorporate Kalman filtering updates but that uses all PSA readings retrospectively to estimate the regression parameters for each patient (all Readings).

Note that, while initially the time of nadir distribution is highly dependent on the prior parameter estimates within each cluster, as new readings become available, the distribution of the time of nadir using Kalman filter updates approaches the distribution of  $tnadir_i$ .

**Figure 3.11. Distribution of the estimated time of nadir among patients after each patient has had 1, 2, 3 or 4 PSA readings (Reading 1, Reading 2, Reading 3 and Reading 4, respectively).**




Readings are taken every 2 months on average. The estimated time of nadir is obtained by recursively applying the Kalman filter and Equation (3.10) for each patient 1, 2, 3 or 4 times. For comparison purposes, we include the distribution of the estimated time of nadir  $tnadir_i$  based on  $\hat{\theta}_i$  (all Readings).  $\hat{\theta}_i$  is obtained by fitting Equation (3.1) retrospectively for each patient given all PSA readings from the start of the hormone treatment to the start of the RT.

We continue to update the probability a patient is in each cluster using Equation (3.29) by comparing the projected PSA, based on the patient's trend in each cluster, to the observed PSA. Patients are assigned to the cluster with highest probability. Figure 3.12 summarizes how clustering classification changes over time.

As the patient's curve parameters are updated, Clusters 1 and 2 become very similar, making it increasingly difficult to discern between the clusters. Also, as additional readings become available, fewer patients are incorrectly classified as being in Cluster 3. This is important to clinicians, since patients in the third cluster tend to need additional medication, which is costly and has associated side effects.

**Figure 3.12. Change of cluster classification based on the number of times the probability of being in each cluster is updated (1, 2, 3 or 4 times, represented as Reading 1, Reading 2, Reading 3, and Reading 4).**



Cluster classification is updated using Equation (3.29). PSA readings are taken every 2 months on average.

### 3.2.2 Clinical Decision Making

We now illustrate the capabilities of our model by comparing two heuristic decision rules for starting RT based on the previous model to the current clinical guidelines. Each month is assumed to consist of 30 days. Since we cannot observe the true time of nadir for a patient, we focus on the patient's estimated time of nadir for comparison. The estimated time of nadir for each patient is obtained from Equation (3.10), based on the regression parameters calculated retrospectively given all PSA readings for that patient. With the exception of the current protocol, all policies use



the proposed Kalman filter model to update the parameters of the PSA curve as new PSA values become available:

Current protocol: In British Columbia, patients who receive neoadjuvant hormone therapy start their radiotherapy treatment if at least one of the following conditions occurs (BC Cancer Agency 2009):

- 8 months of hormone therapy have been received
- PSA levels start to rise
- $\text{PSA} < 0.05 \text{ ng/ml}$  after 4 months

Cumulative probability policy (based on our model): We focus on the cumulative probability of having reached the nadir. Two cumulative probability policies have been analyzed.

- **Cumulative probability policy A**: This policy starts the radiotherapy treatment of the patient if the cumulative probability of having reached the nadir, from the time the hormone therapy started until the time of the latest PSA reading, is greater than a threshold.
- **Cumulative probability policy B**: This policy starts the radiotherapy treatment of the patient if the cumulative probability of having reached the nadir, from the time the hormone therapy started until two months after the latest PSA reading, is greater than a threshold. This policy indicates that RT should be started on a given patient if the probability of having reached the nadir - or of reaching it before the next PSA reading - is greater than a given threshold.

Possible thresholds explored varied from 65% to 90%. If the given threshold is not reached by the eighth month, it is assumed that all patients receive RT at the eighth month.

Threshold probability policy (based on our model): This family of policies focuses on the probability of reaching the nadir in the current period. Four threshold probability policies have been analyzed:

- **Threshold probability policy A**: This policy starts the radiotherapy treatment of the patient if the patient's probability of reaching the nadir from the time of the PSA reading until the

next PSA reading is taken (assumed to be two months afterwards) is higher than a given threshold.

- **Threshold probability policy B:** This policy starts the radiotherapy treatment of the patient if the patient's probability of reaching the nadir from *one month before* the PSA reading until *one month after* that reading is higher than a given threshold.
- **Threshold probability policy C:** This policy starts the radiotherapy treatment of the patient if the patient's probability of reaching the nadir from the time of the current PSA reading until a month after that reading is higher than a given threshold.
- **Threshold probability policy D:** This policy starts the radiotherapy treatment of the patient if the patient's probability of reaching the nadir from *two weeks before* the PSA reading until *two weeks after* that reading is higher than a given threshold.

We considered a variety of possible thresholds between 15% and 60%. If the given threshold is not reached by the eighth month, it is assumed that all patients receive RT at the eighth month.


We used our baseline approach to model the policies. This approach:

- Keeps track of the curve parameters for the three clusters.
- Updates the parameter estimates and the probability of being in each cluster based on the PSA values observed.
- Assigns each patient to the cluster with the highest probability  $P(j|i,t)$  in each period.
- Uses the integral approximation (3.11) to estimate the nadir probabilities given that the patient belongs to the cluster with highest probability.
- Uses equation (3.27) to weight the prior estimates of the parameters within each cluster.
- Does not incorporate measurement error.

Policies are compared based on the mean absolute difference between the policy's treatment time and the patient's estimated time of nadir. From Figure 3.13 we see that, among the policies explored, 85% cumulative probability policy B and 15% threshold probability policy D perform closest to the estimated time of nadir (lowest mean absolute difference). Their mean absolute difference from the time of nadir is 29 days and 36 days, respectively, compared to 45 days under

the current protocol. Additional analysis is required to determine the patient specific threshold that achieves the smallest difference from the estimated time of nadir.

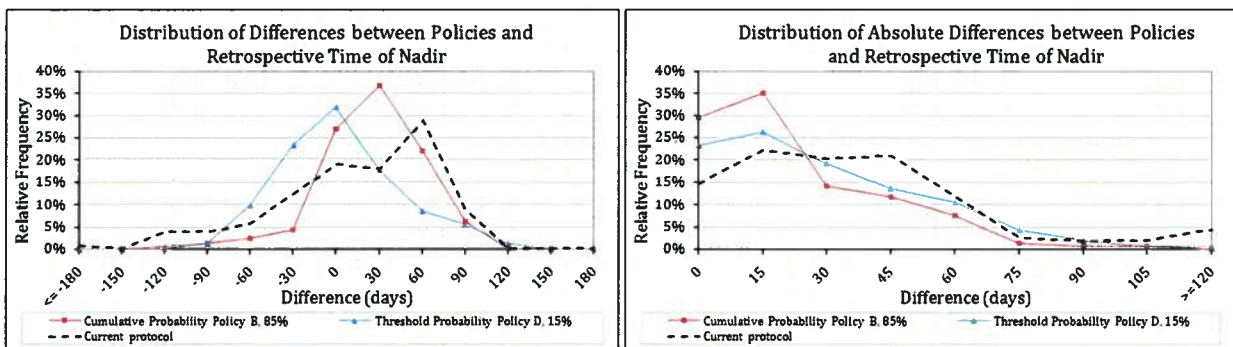
**Figure 3.13. Comparison of the mean absolute difference between the RT start time based on each policy and the estimated time of nadir for each patient in the population.**



The estimated time of nadir is based on Equations (3.1) and (3.10) using all PSA readings for that patient (not achievable). Within each policy, various thresholds are compared. The policies are compared to the current protocol (dashed line).

The distribution of the differences and absolute differences of the 85% cumulative probability policy B and 15% threshold probability policy D are compared to the current protocol in Figure 3.14.

**Figure 3.14. Comparison of the population distribution of the difference and absolute difference of the estimated time of nadir for each patient to the RT start time of that patient based on the 85% cumulative probability policy B, 15% threshold probability policy D and the current protocol.**



The estimation of the time of nadir for each patient is based on Equations (3.1) and (3.10) using all PSA readings for that patient (not achievable).

Not only does the current protocol have a larger mean difference from the estimated time of nadir, but it also has greater variability (with a variance of 1024, compared to variances of 431 and 612 for the 85% cumulative probability policy B and the 15% threshold probability policy D, respectively). Under the current protocol, most patients will wait for a rise in their PSA or for eight months to have elapsed to start RT. However, based on our analysis, over 82% of the patients will have reached their nadir by the seventh month. If starting the RT treatment at the PSA nadir is assumed to be best for patients, by following the current protocol, the RT treatment is started a month and a half later, on average, than the ideal time to initiate treatment. Starting the RT treatment too late could, in theory, be linked to disease progression, increased risk of cells becoming resistant to treatment, and greater toxicity associated with hormone therapy.

We have applied the 85% cumulative probability policy B and the 15% threshold probability policy D based on parameter priors from this dataset to another dataset. Under the 85% cumulative probability policy B, 75% of the patients would have started RT earlier than what their current protocol suggested. Under the 15% threshold probability policy D, the number of patients that would have started earlier than their current protocol was 89%. The average difference between the time of treatment and the time of treatment under the current protocol is 74 and 103 days, respectively.

We then considered the following modifications to the baseline model:

Simulation: This modeling approach calculates the probability of the nadir by using simulation instead of the integral approximation described in equation (3.11). Simulation results are obtained based on 10,000 replications.

Weight cluster probabilities: Rather than assigning each patient to the cluster with the highest probability  $P(j)_{i,t}$  in each period, this approach weights the nadir probability within each cluster times the probability  $P(j)_{i,t}$  that the patient belongs to the cluster.

Weight priors: This modeling approach weights the regression coefficients by their variance to obtain the prior estimates within each cluster. Equation (3.28) is used instead of equation (3.27).

Include measurement error: This approach incorporates the covariance matrix of the initial parameters as measurement error.

As can be seen in Table 3.3, none of the modifications to our modeling approach improved our solutions significantly, suggesting that our initial assumptions are appropriate.

**Table 3.3. Model comparisons based on the mean absolute difference and maximum absolute difference between each policy's RT start time and the estimated *tnadir* for all patients.**

Range of probability	Threshold Used	Initial		Simulation		Weight probabilities.		Use equation (16)		Include msmt. error		No clusters	
		mean	max	mean	max	mean	max	mean	max	mean	max	mean	max
Threshold Probability Policy A	0.15	60	154	59	154	60	228	60	228	62	174	66	228
	0.2	58	154	58	154	59	228	59	228	59	154	61	228
	0.3	56	145	56	145	56	145	57	228	57	145	56	228
	0.4	55	145	55	145	52	145	52	228	54	145	49	145
	0.5	49	133	50	137	47	130	50	228	49	137	49	137
	0.6	47	130	46	130	48	130	49	228	49	130	49	130
Threshold Probability Policy B	0.15	51	145	51	145	50	228	50	228	53	145	54	228
	0.2	49	145	49	145	48	228	43	228	50	145	51	228
	0.3	39	130	40	130	37	127	36	228	44	133	41	228
	0.4	36	121	36	121	35	119	36	228	41	127	40	119
	0.5	40	127	40	127	43	127	42	228	42	119	44	127
	0.6	46	127	46	127	47	127	48	228	46	127	48	127
Threshold Probability Policy C	0.15	52	145	52	145	31	119	51	228	54	145	52	228
	0.2	49	145	50	145	30	119	46	228	49	133	50	228
	0.3	43	127	43	127	29	106	50	228	45	130	45	130
	0.4	47	127	46	127	29	106	49	228	47	127	47	127
	0.5	49	127	49	127	31	106	49	127	49	127	48	127
	0.6	49	127	49	127	31	106	49	127	49	127	49	127
Threshold Probability Policy D	0.15	36	121	37	127	34	121	36	228	41	133	40	228
	0.2	36	121	36	121	37	119	38	228	40	121	42	228
	0.3	45	127	45	127	46	127	47	228	43	119	46	127
	0.4	48	127	48	127	49	127	49	127	47	127	49	127
	0.5	49	127	49	127	49	127	49	127	49	127	49	127
	0.6	49	127	49	127	49	127	49	127	49	127	49	127
Cumulative Probability Policy A	0.65	32	106	31	106	33	106	36	106	32	106	32	106
	0.7	34	106	34	106	36	106	38	106	33	106	34	106
	0.75	36	106	36	106	39	106	41	106	35	106	38	106
	0.8	39	106	39	106	41	106	45	106	39	106	41	106
	0.85	43	106	43	106	44	106	47	106	43	106	43	106
	0.9	46	106	46	106	46	95	48	106	46	106	47	106
Cumulative Probability Policy B	0.65	37	130	38	130	35	130	32	228	37	130	35	130
	0.7	33	119	33	119	34	119	30	228	33	119	34	130
	0.75	30	119	30	119	31	119	31	228	33	119	32	119
	0.8	29	119	29	119	31	106	30	228	29	106	33	119
	0.85	29	106	29	106	32	106	31	106	32	106	32	106
	0.9	32	106	32	106	35	106	38	106	33	106	34	106

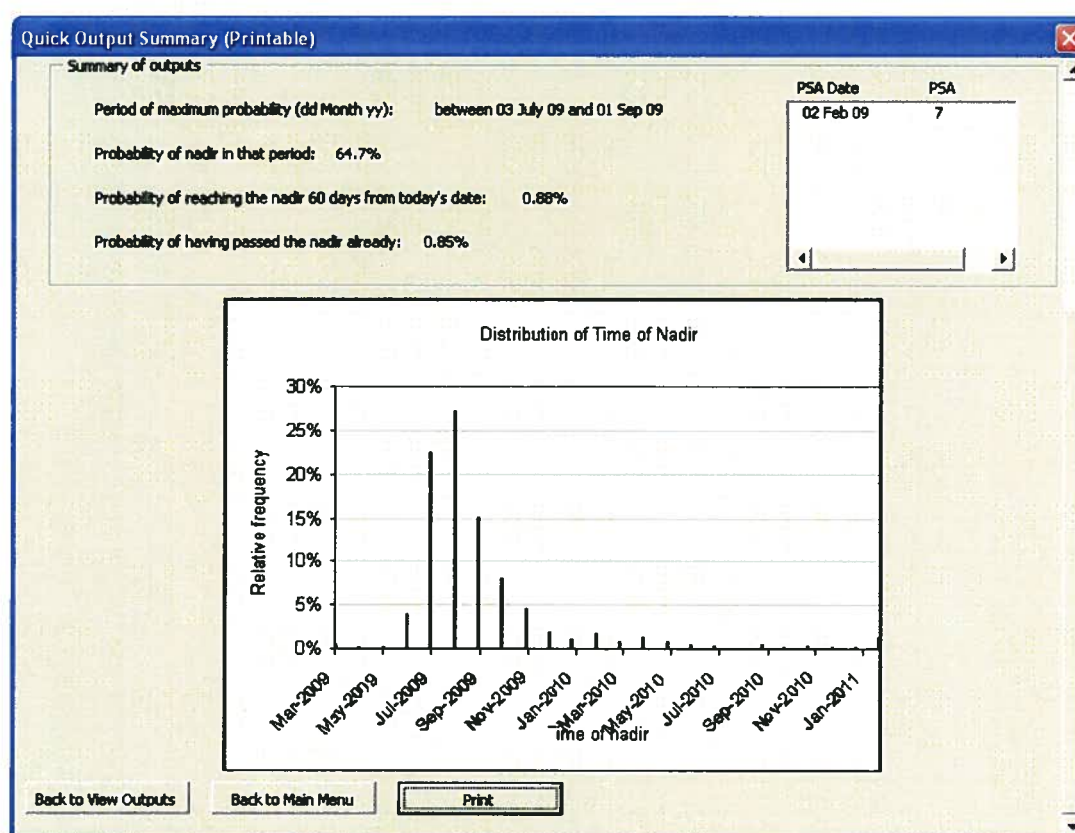
The minimum mean and maximum difference within each policy and modeling approach are highlighted.

In addition, we compared our results to the results obtained if clustering was omitted from our analysis. That is, we assumed that all patients in the population belong to a single cluster (no clusters). While the absolute difference between the treatment time and the retrospectively estimated time of nadir under all policies without clustering is not smaller than if clustering was incorporated, the difference between both models is not large enough to warrant the additional

complications involved in keeping the clusters. Note, however, that in a more heterogeneous population, clustering might have a larger impact on the solution.

To implement this model, we have developed a user friendly Excel based tool that can be run at the clinician's computer. A screenshot of the output page of the tool appears in Figure 3.15. The main inputs of this tool include the hormone start date and the PSA dates and values. The tool then provides clinicians with outputs related to the estimated distribution of the time of nadir. Clinicians can use this tool to decide whether to start RT treatment based on the probability at any given period of treating at the nadir as well as to measure tradeoffs of delaying the RT treatment. The model also provides the physician, and the patient, with an estimate of when the patient's RT is likely to start several months in advance. As additional PSA values become available, they can be easily incorporated by using the tool, and the outputs are recalculated accordingly. The tool interacts with an Access database to record all inputs and outputs obtained.

**Figure 3.15. Sample output obtained from the tool**



Given that the PSA dynamics might depend on the type of hormone treatment, before this tool can be used prospectively, clinicians need to calibrate the model. That is, clinicians should first apply the model retrospectively to their patients. The tool would be used to compare the model

estimates to their own assessment of the time of nadir. If the model estimates are not accurate, it is necessary to revise the prior distribution of the parameters based on the PSA dynamics of that group of patients. Additional validation after the model has been calibrated is suggested. After the model has been calibrated and validated, clinicians might use the tool to make prospective decisions of patients with similar characteristics to the patients used in the validation and calibration process.

### 3.3 Conclusions and Remarks

We have developed an iterative approach to update the estimates of the distribution of the PSA nadir of prostate cancer patients receiving neoadjuvant hormone therapy treatment prior to radiation therapy. By using a threshold to decide whether to start RT, we are able to identify earlier when the nadir is likely to be reached. Furthermore, we are able to decrease the variability of the difference between the RT treatment time and the estimated time of nadir. Additional analysis is needed to determine the optimum threshold and the patient specific policy that maximizes the probability of treating at the nadir.

To determine the best patient specific policy, this problem might be formulated as a discrete time, finite horizon, Markov decision problem (MDP). The objective of such a model is to maximize the probability of treating at the nadir. The decision epochs correspond to the times at which PSA readings are taken. The two possible actions are to start RT or to wait for another PSA reading. If RT is started, the patient receives a reward based on the probability of reaching the nadir in the following period, and no additional PSA readings are taken. If it is decided to wait, no reward is received. In that case, a new PSA reading is taken in the following period, and the process is repeated. The state of the model is the patient's PSA curve parameters and the covariance of the parameters used to calculate the distribution of the time of nadir. The state is updated using the Kalman filter.

Note that, while this model entails a continuous, partially observable state space, given the Kalman filter properties discussed, the parameter covariances do not depend on the new PSA readings and can be computed off-line based on the prior information. In addition, while the reward function is based on the ratio of correlated normal random variables, the integral approximation discussed in this paper would simplify the reward calculations.

This model assumes that there is a fixed time at which readings are taken, and that RT will be started at the last period if it has not started before. Other questions to be addressed include when to take the next reading and what a good endpoint for the model might be. While eight months have been assumed throughout the paper, it is possible that a longer planning horizon might be appropriate.

A key assumption of our model is that the PSA nadir is directly linked to better patient outcomes. The validation step is to design a formal clinical trial comparing a fixed time interval, or a target PSA, to the estimate of the nadir based on our decision model. Such a trial would compare the proposed model to the current protocol in terms of survival and time to PSA relapse.

Once this model has been validated, clinicians could use it prospectively to decide when to start RT on their patients. Using the user friendly tool developed, clinicians would input the PSA values of their patients in each period. They would then obtain the estimated probability of reaching the nadir within 15 days and the cumulative probability of having reached the nadir or of reaching it within the next period. Depending on the policy chosen, the patient would start RT if the threshold probability or the cumulative probability is above 15% or 85%, respectively.



## CHAPTER 4: VALIDATION OF A PSA DYNAMICS MODEL AND NADIR PREDICTION METHODS

In the previous chapter, we developed a novel approach to modeling the individual disease progression of prostate cancer patients, based on the dynamics of their prostate specific antigen (PSA). In the model, estimates are updated as new patient-specific information becomes available. The model was developed for a cohort of intermediate and high risk prostate cancer patients who were enrolled in a prospective randomized trial (Morris et al. 2009). Under the assumption that it is best to treat patients when their PSA is as low as possible (Gleave et al. 2000), we used our model to decide when to initiate radiation therapy based on when the PSA reached its nadir level.

We now proceed to validate our model by applying it to a larger, more heterogeneous cohort of patients from the British Columbia Cancer Agency (BCCA) database. We avoid bias by only using the models and policies developed from the original cohort of patients to estimate the best time of treatment for this new cohort of patients. We show that, even though the models were not calibrated to the more heterogeneous validation dataset, we are better able to plan optimal treatment closer to the patients' PSA nadir in comparison with the current protocol (time of RT). Section 4.1 summarizes the methodology used. Section 4.2 describes the data used in the validation process. The model is validated in Section 4.3. We then test the policy implications of using our modeling approach in Section 4.4. We conclude with final remarks in Section 4.5.

### 4.1 Methodology

In order to assess our model we must validate:

- PSA curve dynamics
- Update the curve as new PSA readings become available
- Cluster classification
- Policy implications of our model

Key to our validation process is having an estimate of when each patient reached his minimum PSA level. Two alternatives are:

- The retrospective nadir estimated by fitting a curve through all PSA values for each patient.

- The time - or an interval around the time - at which the minimum PSA value for each patient was observed.

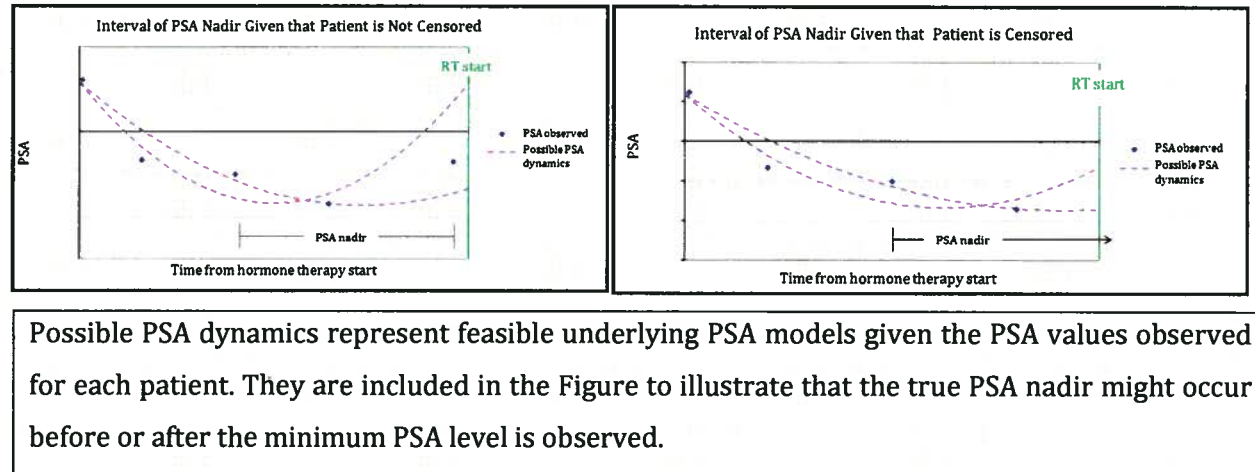
While the second option avoids any modeling assumptions, it is highly dependent on how often the PSA readings were taken and the time at which radiation therapy (RT) was started. Patients' PSAs might be censored, that is, their minimum observed PSA level might correspond to the last reading obtained before RT started. For those patients, however, their true PSA nadir might have been reached later. This is not fully accommodated if we only examine the minimum observed PSA value before the start of RT. On the other hand, the retrospective nadir estimated from a model fit to all the data would not depend on the frequency of the PSA readings and might be obtained by extrapolating the model parameters after RT is started. While we acknowledge that in an ideal scenario, such extrapolation would be avoided, it might provide a more reliable nadir estimate for censored PSA readings. We use both approaches to estimating the nadir in our validation process.

We begin by testing the goodness of fit of our PSA dynamics curve. For that purpose, we fit the proposed curve based on all observations for each patient in the database. For those patients that have at least 4 PSA observations (over 50% of the patients analyzed), we compute the R squared values of their curve. MSE are also calculated to obtain an idea of the magnitude of the residuals.

Next, we compute the observed PSA nadir and compare it to the estimated PSA nadir based on each patient's curve parameters. We first determine the time at which the minimum PSA is first observed. It is assumed that there is a strong likelihood that the true PSA nadir occurs after the previous PSA reading and before the following PSA reading. An alternative to this approach would be to select a fixed interval around the PSA nadir within which the nadir is assumed to occur. While selecting a fixed interval would avoid the dependency on how often the PSA readings were taken, a fixed interval would depend instead on what the chosen length of the interval is. Further analysis is necessary to determine such an interval.

If the minimum PSA observed occurs right before RT starts (a censored patient), it is only assumed that there is a high likelihood that the observed PSA nadir occurs after the previous PSA reading. Figure 4.1 illustrates the intervals assumed to have high likelihood of containing the PSA nadir for two sample patients.

**Figure 4.1. Illustration of interval of PSA nadir based on whether patient is censored.**



The estimated PSA nadir is calculated from Equation (3.10). To assess how well such estimates perform, we calculate the proportion of patients for whom the estimated PSA nadir falls within the interval assumed to contain the patient's PSA nadir.

The curve parameters are updated as new PSA readings become available. Using the prior distribution of the parameters obtained from our original dataset, we apply the Kalman filter (baseline) model described in Chapter 3 to our validation dataset. Each time the curve parameters are updated, we estimate the nadir (from Equation 3.10) and compute the proportion of patients for which the estimated nadir falls within the interval assumed to contain the observed PSA nadir. Moreover, we compute the distribution of the difference between the estimated nadir obtained from the updated curve parameters and the estimated nadir obtained retrospectively by fitting Equation 3.1 based on all PSA readings for each patient.

We then assess our cluster classification. Given the patients' initial PSA levels, they are classified into the three clusters obtained from the original clinical trial. We estimate the proportion of patients in each of the original clusters and how cluster classification changes over time. The PSA nadir estimations given the original cluster classification are compared to the nadir estimations under no clustering assumptions.

Finally, we study the policy implications of our model. We apply the 15% threshold probability policy and the 85% cumulative probability policy to the new cohort of patients and assess how the time of treatment under both policies and under the current protocol differ from the estimated time of nadir. Moreover, we show that under the proposed policies patients would have been

treated closer to the time at which the minimum PSA value was observed than the actual time at which they received RT (current protocol).


#### 4.2 Data Description

Our database consists of 422 patients who received up to a year of hormone therapy before curative intent radiation therapy and had at least 3 PSA readings (and up to 11 PSA readings) throughout their hormone treatment. Over 50% of the patients had at least 4 PSA readings throughout their hormone treatment. All cases were diagnosed between January 2000 and May 2007. None of the patients participated in the clinical trial used to develop the model. The initial PSA - used as the first PSA reading in our model - is the last reading taken before hormone therapy started. The reading might also have been taken on the day hormone therapy started. Patients received either luteinizing hormone-releasing (LHRH) therapy - with only the first month of LHRH covered by non-steroidal anti-androgen (NSAA) - or total androgen blockade (with both LHRH and NSAA throughout the course of hormone therapy). Patients are classified as:

- **Low risk patients (LR):** Includes patients with all:
  - Gleason  $\leq 6$
  - Initial PSA  $\leq 10$
  - Histologically-proven prostate cancer stage T1a-T2b (AJCC 2002)
- **Intermediate risk patients (IR):** Includes patients with histologically-proven prostate cancer stage T1a-T2c and either:
  - Gleason  $\leq 6$  and  $10 < \text{initial PSA} \leq 20$
  - Gleason = 7 and initial PSA  $\leq 20$
  - Tstage T2c (AJCC 2002), Gleason  $< 8$ , and initial PSA  $\leq 20$
- **High risk patients (HR):** Includes patients with either:
  - Gleason  $\geq 8$
  - Initial PSA  $> 20$
  - Histologically-proven prostate cancer stage T3a-T4 (AJCC 2002)

Our validation database is described in Figure 4.2 and Table 4.1.

**Figure 4.2. Proportion of patients in the validation database.**



**Table 4.1. Number of patients in the validation database.**

Eligibility for initial clinical trial	Hormone therapy	Risk group			Total
		LR	IR	HR	
Eligible	LHRH monotherapy with only first month of NSAA	0	71	157	228
	Total androgen blockade until RT	0	19	48	67
Not eligible	LHRH monotherapy with only first month of NSAA	5	0	85	90
	Total androgen blockade until RT	3	0	34	37
<b>Total</b>		8	90	324	422

While only intermediate and high risk non-metastatic prostate cancer patients with an initial PSA below 40 ng/ml and histologically-proven prostate cancer stage T1c -T3a (AJCC 2002) would have been eligible to participate in the clinical trial (Morris et al. 2009) used to develop the original model, our validation database incorporates a wider group of patients. From Figure 4.2 and Table 4.1, we see that 295 patients - or 69.9% of the patients - would have been eligible to participate in the clinical trial used to develop the initial model. Most of the non eligible patients would have been excluded because of their high initial PSA or their low risk classification. Moreover, 25% of the patients received total androgen blockade, and most patients were either high risk or intermediate risk patients.

On average, patients received 216 days of hormone therapy prior to radiation therapy (see Table 4.2). For all risk groups, patients that received total androgen blockade received slightly longer hormone therapy compared with patients that received LHRH monotherapy.

Table 4.2. Length of hormone treatment for patients in the validation database (Average Number of Days [Minimum, Maximum]).

Eligibility for initial clinical trial	Hormone therapy	Risk group		
		LR	IR	HR
Eligible	LHRH monotherapy with only first month of NSAA	0	214 [70,338]	215 [32,345]
	Total androgen blockade until RT	0	240 [148,342]	228 [68,376]
Not eligible	LHRH monotherapy with only first month of NSAA	213 [183,268]	0	207 [63,341]
	Total androgen blockade until RT	261 [216,323]	0	214 [88,358]

PSA readings were taken every 68 days, on average. High risk patients had PSA readings completed more often than did intermediate and low risk patients (see Table 4.3). Within the group of high risk patients, patients that would not have been eligible to participate in the initial clinical trial had PSA readings completed more often than did patients that would have been eligible.

Table 4.3. Time (in days) between PSA readings for patients in the validation database (Average, [Minimum, Maximum]).

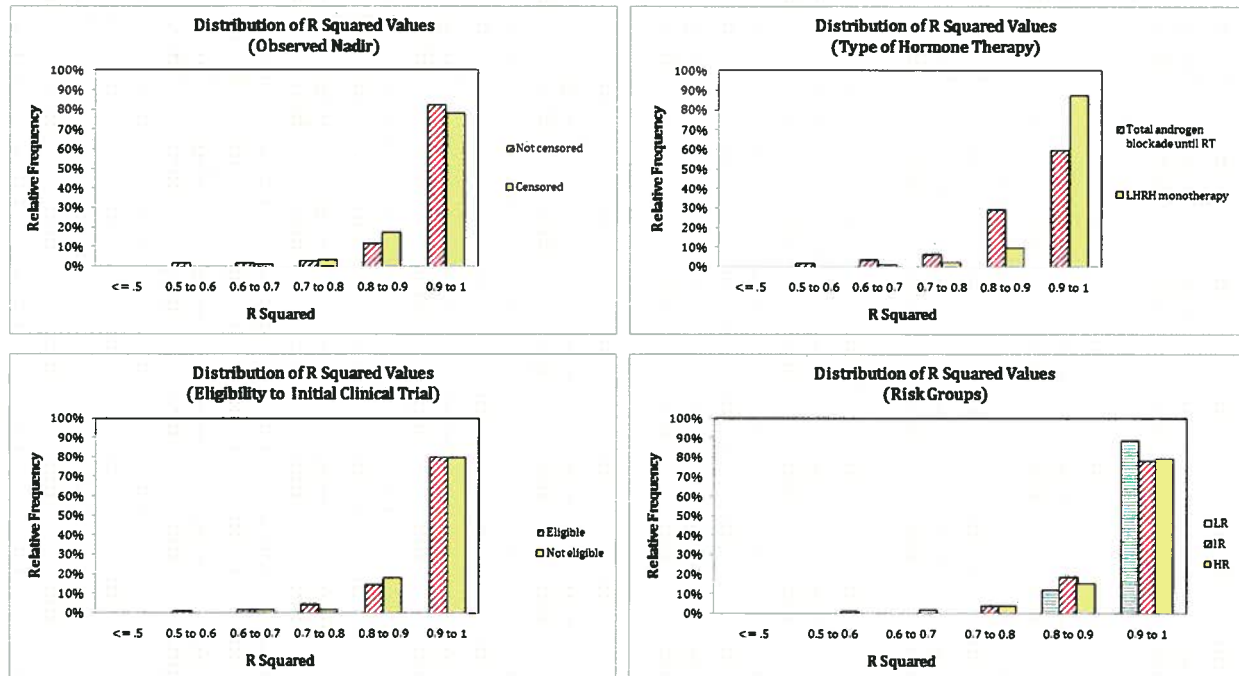
Eligibility for initial clinical trial	Hormone therapy	Risk group		
		LR	IR	HR
Eligible	LHRH monotherapy with only first month of NSAA	0	70 [20,140]	69 [11,130]
	Total androgen blockade until RT	0	76 [40,158]	70 [20,127]
Not eligible	LHRH monotherapy with only first month of NSAA	79 [55,102]	0	66 [20,125]
	Total androgen blockade until RT	92 [84,98]	0	61 [17,122]

### 4.3 Model Validation

As mentioned in Chapter 3, we assume that the PSA progression of patient  $i$  as a function of the time  $t$  elapsed since the patient started his hormone treatment (in days) can be modeled by a log quadratic curve (3.1):  $Y_{i,t} = \ln(PSA_{i,t}) = \alpha_i + \beta_i t + \gamma_i t^2 + \varepsilon_i$ ,  $\varepsilon_i \sim N(0, V_i)$ .

We proceed to fit (3.1) to all patients in our validation database based on all readings for each patient. Figure 4.3 summarizes the  $R^2$  obtained.

**Figure 4.3. Distribution of  $R^2$  values for patients that had at least four PSA readings.**



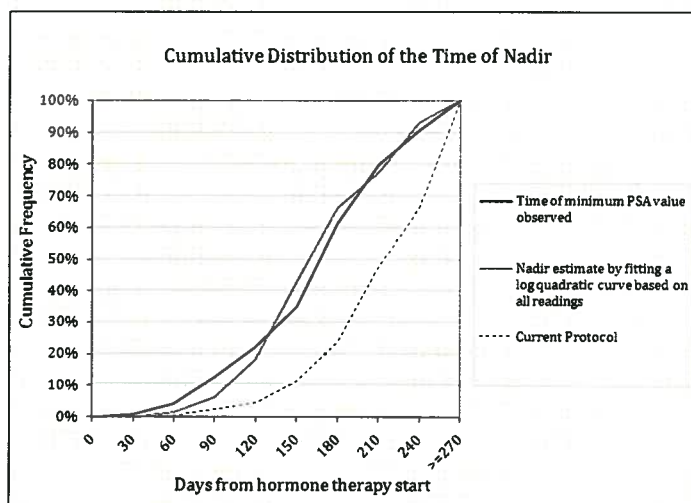
As can be seen in Figure 4.3, the log quadratic curve seems to fit the PSA dynamics of the patients well. For those patients with at least 4 observations - and an average of 5 observations - the  $R$  squared values ranged from 0.56 to 0.99. While we acknowledge that, given the limited number of PSA readings available for each patient, the  $R$  squared values might not be a robust indicator of the goodness of fit of the curve, high  $R$  squared values provide some evidence that the log quadratic curve fits the available data well. The MSE of the curves ranged from .00015 to 3.19 with an average MSE of 0.34.

Whether patients would have been eligible to participate in the original trial had no impact on how well the quadratic curve fit the PSA dynamics. Note, however, that the curve tends to better fit patients that had LHRH monotherapy with only the first month of NSAA.

A comparison of the cumulative distribution of the retrospective time of nadir obtained by fitting a log quadratic curve through all PSA readings for each patient versus the distribution of the time at which the minimum PSA value for each patient was observed is presented in Figure 4.4. Note that, while some differences are to be expected given that our second measure is highly dependent on

when the PSA readings were taken, both distributions are similar, providing further evidence that the log quadratic curve is appropriate.

**Figure 4.4. Comparison of the cumulative distribution of the time of the minimum PSA value observed versus the nadir estimated by fitting a log quadratic curve based on all readings**



Only 25% of the patients in our validation database were not censored (see Table 4.4): the PSA levels rose or did not drop further after reaching their minimum observed level. Censored patients may have been treated too early or the PSA may not have been measured frequently enough to observe the PSA relapse.

**Table 4.4. Proportion of not censored patients in the validation database.**

Eligibility for initial clinical trial	Hormone therapy	Risk group		
		LR	IR	HR
Eligible	LHRH monotherapy with only first month of NSAA	0%	15%	30%
	Total androgen blockade until RT	0%	26%	25%
Not eligible	LHRH monotherapy with only first month of NSAA	20%	0%	19%
	Total androgen blockade until RT	0%	0%	44%


Note that the assay sensitivity ranges from 0.026 ng/mL to 0.052 ng/mL (Niblock et al 2006). Given assay sensitivity and other patient specific fluctuations, it is possible that subsequent PSA



readings for not censored patients might be lower - that not censored patients are in fact censored. Conversely, it is also possible that a patient that is assumed to be censored has in fact reached his minimum PSA value before RT started.

We compare the proportion of patients for whom the estimated PSA nadir falls within the interval assumed to contain the patient's PSA nadir to the proportion of patients for whom their time of RT (current protocol) falls within the interval assumed to contain the patient's PSA nadir (see Figure 4.5 and Table 4.5).

**Figure 4.5. Comparison of the estimated PSA nadirs to the current protocol based on whether they fall within the interval assumed to contain the true nadir.**



Two possible nadirs are estimated for each patient: under no additional assumptions and under the assumption that RT must start by the eighth month as suggested by current guidelines. The current protocol represents the actual time each patient received their radiotherapy treatment. Note that, given our definition of the interval of PSA nadir, the current protocol will always fall within that interval if patients are censored and will not fall within the interval if patients are not censored. We incorporate the current protocol given other patient characteristics to illustrate possible gains obtained by applying our model to estimate the nadir. For the current protocol, the proportion of patients that do not fall within the interval of PSA nadir can be interpreted as the proportion of patients for whom RT treatment was started after the interval of PSA nadir.

**Table 4.5. Proportion of patients for whom the estimated nadirs and the current protocol fall within the interval assumed to contain the PSA nadir**

		Estimated nadir	Estimated nadir (limit maximum neoadjuvant hormone duration to eight months)	Current protocol
Observed nadir	Not censored	91%	92%	0%
	Censored	93%	92%	100%
Hormone therapy	Total androgen blockade	93%	91%	69%
	LHRH monotherapy	92%	92%	76%
Eligibility for initial clinical trial	Eligible	91%	91%	75%
	Not eligible	95%	95%	75%
Risk group	LR	100%	100%	88%
	IR	92%	92%	82%
	HR	92%	92%	72%
Total	All Patients	92%	92%	75%

Whether we assume that RT should be started at the eighth month, at the latest, does not have a significant impact on the proportion of patients for whom the estimated PSA nadir falls within the interval assumed to contain the patient's PSA nadir. From Figures 4.4 and 4.5 and Table 4.5, note that the estimated PSA nadir outperforms the current protocol. However, the estimated nadir is based on all observations and can only be obtained retrospectively. We proceed to compare the nadir estimated after the patient has had 1, 2 or 3 readings to the current protocol. Starting from the priors obtained from the original clinical trial - under clustering and no clustering assumptions - we use the proposed Kalman Filter model prospectively to update the distribution of the parameters as new readings become available. Table 4.6 compares the proportion of patients for whom the estimated nadir after one, two or three readings falls within the interval assumed to contain the patient's PSA nadir.

From Table 4.6 we can see how, with just the first PSA reading, the proportion of patients for whom the estimated nadir falls within the interval assumed to contain the patient's PSA nadir is higher than under the current protocol. Clinicians might use this information to give patients an initial estimate of when their RT treatment will be given, as well as to plan the RT resources accordingly. Given how wide the interval is - ranging from 21 days to over 8 months - and the high

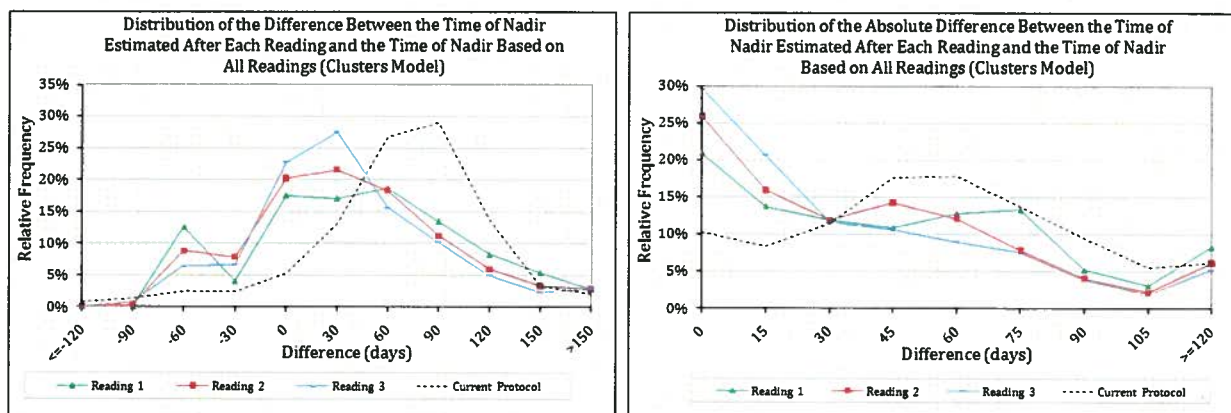
percentage of censored patients, this proportion only increases slightly as new readings become available. Intervals for non-censored patients are narrower than for censored patients.

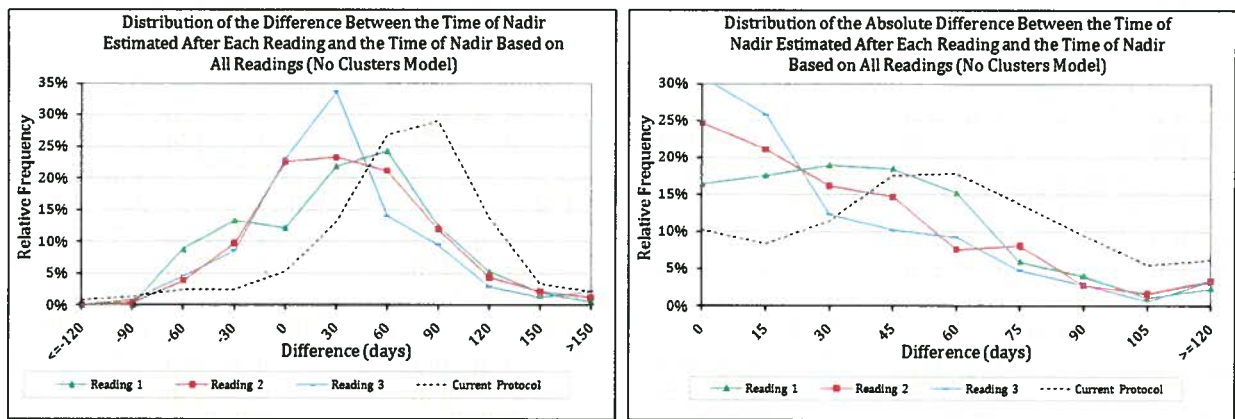
**Table 4.6. Comparison of the estimated PSA nadir after the patient has had one, two or three readings to the current protocol based on the proportion that falls within the interval assumed to contain the patient's PSA nadir.**

		Reading 1		Reading 2		Reading 3		Current protocol
		No Cluster	Cluster	No Cluster	Cluster	No Cluster	Cluster	
Time of nadir	Not censored	46%	45%	43%	48%	56%	55%	0%
	Censored	91%	93%	92%	90%	90%	91%	100%
Hormone therapy	Total androgen blockade	66%	69%	69%	70%	71%	73%	69%
	LHRH monotherapy	84%	84%	83%	83%	84%	85%	76%
Eligibility for initial clinical trial	Eligible	80%	81%	79%	80%	82%	83%	75%
	Not eligible	79%	80%	80%	80%	79%	80%	74%
Risk group	LR	100%	88%	100%	100%	100%	100%	88%
	IR	86%	88%	87%	83%	89%	89%	82%
	HR	77%	78%	77%	78%	78%	79%	72%
Total	All Patients	80%	81%	79%	80%	81%	82%	75%

As can be seen in Figure 4.6, the accuracy of our estimation improves as new PSA readings become available.

**Figure 4.6. Distribution of the difference and absolute difference between the estimated time of nadir based on the Kalman Filter model and the nadir estimated from a model fit to all data.**





The estimated time of nadir is calculated under clustering (clusters model) and no clustering (no clusters model) assumptions after each patient has had one, two or three PSA readings (Reading 1, Reading 2 and Reading 3 respectively).

Note that comparisons in this figure are made under the assumption that the estimated nadir based on all observations is an accurate representation of the true time of nadir.

On average, patients had their second reading 65.6 days after they started hormone therapy and their third reading 138 days after starting hormone therapy. Our results suggest that taking PSA readings more often might improve the accuracy of the PSA nadir estimations. This is key to physicians, since more accuracy in the estimation of the PSA nadir would allow them to know sooner and with less variability when the PSA nadir will be reached and improve planning of the RT resources.

In our baseline model, patients have been classified into the three clusters obtained from the original clinical trial data. The mean estimated time of nadir and minimum PSA value observed by cluster is presented in Table 4.7. Note that patients in the third cluster continue to have an estimated time of nadir that is later than that of the other two clusters.

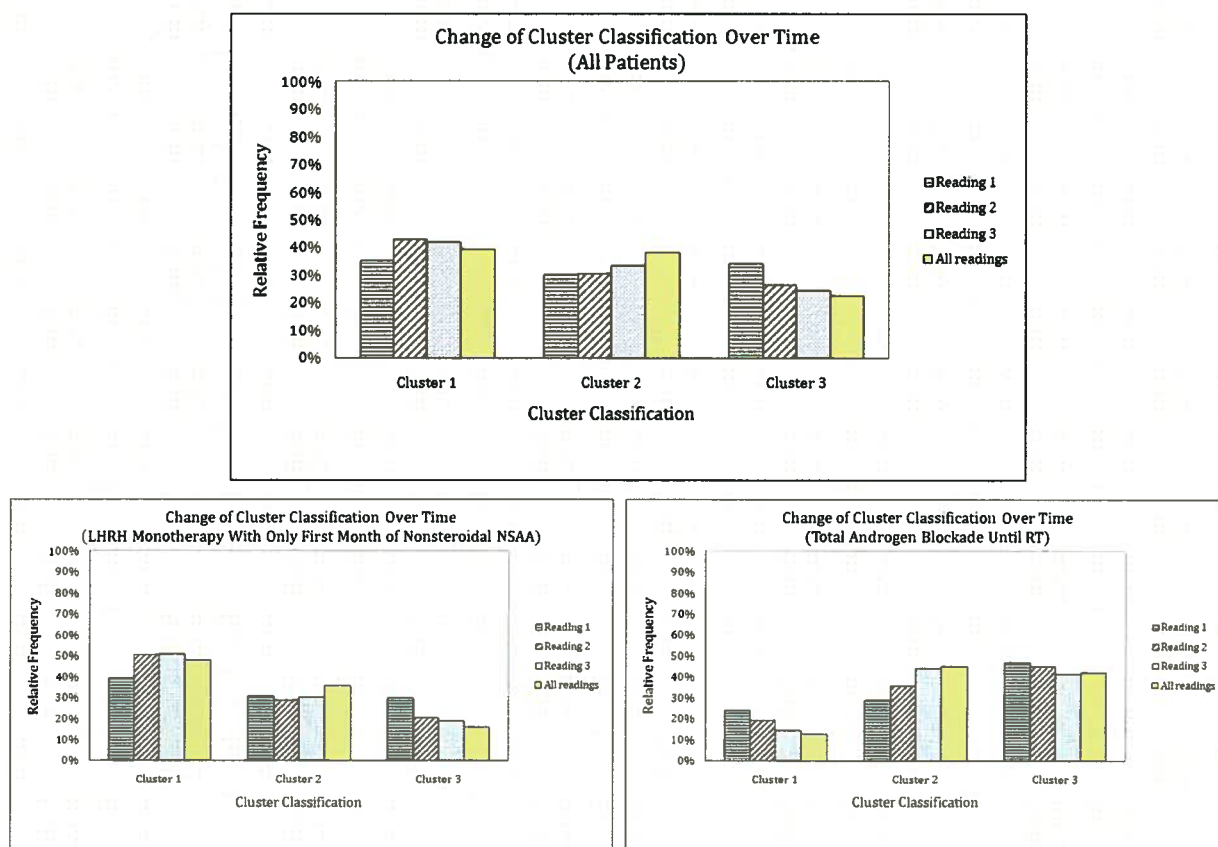
**Table 4.7. Estimated time of nadir and minimum PSA value observed by cluster.**

	Cluster 1	Cluster 2	Cluster 3
Estimated time of nadir (mean, in months)	6.0	5.7	>>8.0
Minimum PSA value observed (mean [variance])	.38[.47]	.8[1.66]	1.4[7.92]

The proportion of patients classified into each of the original clusters is presented in Figure 4.7. Based on our first two readings, 20% of the patients that received LHRH monotherapy are

classified into Cluster 3, while 45% of the patients that received total androgen blockade are classified into the third cluster. While further recalibration of our cluster classification is necessary, our results indicate that patients that have a higher initial PSA value might have a higher likelihood of requiring additional hormones during the treatment.

**Figure 4.7. Change of cluster classification over time for all patients and patients classified based on whether they had LHRH monotherapy or total androgen blockade until RT**




From Figure 4.6 and Table 4.6, note that no significant improvement in the accuracy of the nadir estimation is achieved when patients in our validation database are clustered into three groups based on our original cluster classification. The population used to design the original cluster classification might have been too homogeneous. Patients in this new - more heterogeneous - database should therefore be reclustered, following a similar methodology as highlighted in Chapter 3.

#### 4.4 Policy Implications

As highlighted in Chapter 3, the two policies that performed best on our original database included the 85% cumulative probability policy B and the 15% threshold probability policy D. The 85% cumulative probability policy B (85% policy) starts radiotherapy treatment if the cumulative probability of having reached the nadir - or of reaching it before the next PSA reading is greater than 85%. On the other hand, the 15% threshold probability policy D starts the radiotherapy treatment of the patient if the patient's probability of reaching the nadir from two weeks before the PSA reading until two weeks after that reading is higher than a given threshold. Both policies start the radiotherapy treatment if the given threshold is not reached by the eighth month.

The cumulative distribution of the time of treatment under each protocol is compared to the cumulative distribution of the retrospective nadir estimate in Figure 4.8.

**Figure 4.8. Comparison of the cumulative distribution of the time of treatment under each protocol to the distribution of the nadir estimated based on all readings.**



With the exception of the current protocol, all other estimates are calculated under clustering assumptions. The 15% and 85% probability policies represent the two protocols highlighted in Chapter 3. Reading 2 highlights the distribution of the nadir estimated for each patient after two readings - one reading on or before hormone therapy started and one reading after hormone therapy started. The current protocol represents the actual time patients received RT.



From Figure 4.8, we see that the 15% and the 85% probability policies outperform the current protocol in terms of treating closer to the estimated time of nadir. While treating patients based on the highlighted policies is recommended, the nadir estimated after two readings provide clinicians with an initial idea of when the nadir is expected to be reached.

Next, we compare the two policies under clustering (clusters model) and no clustering assumptions (no clusters model) to the time of RT treatment (current protocol) based on the differences from the estimated time of nadir and from the time of the minimum PSA value observed. The distribution of the differences for all patients is illustrated in Figure 4.9.

**Figure 4.9. Distribution of the difference from the estimated time of nadir and the time of the minimum PSA value observed for the proposed policies and the current protocol.**

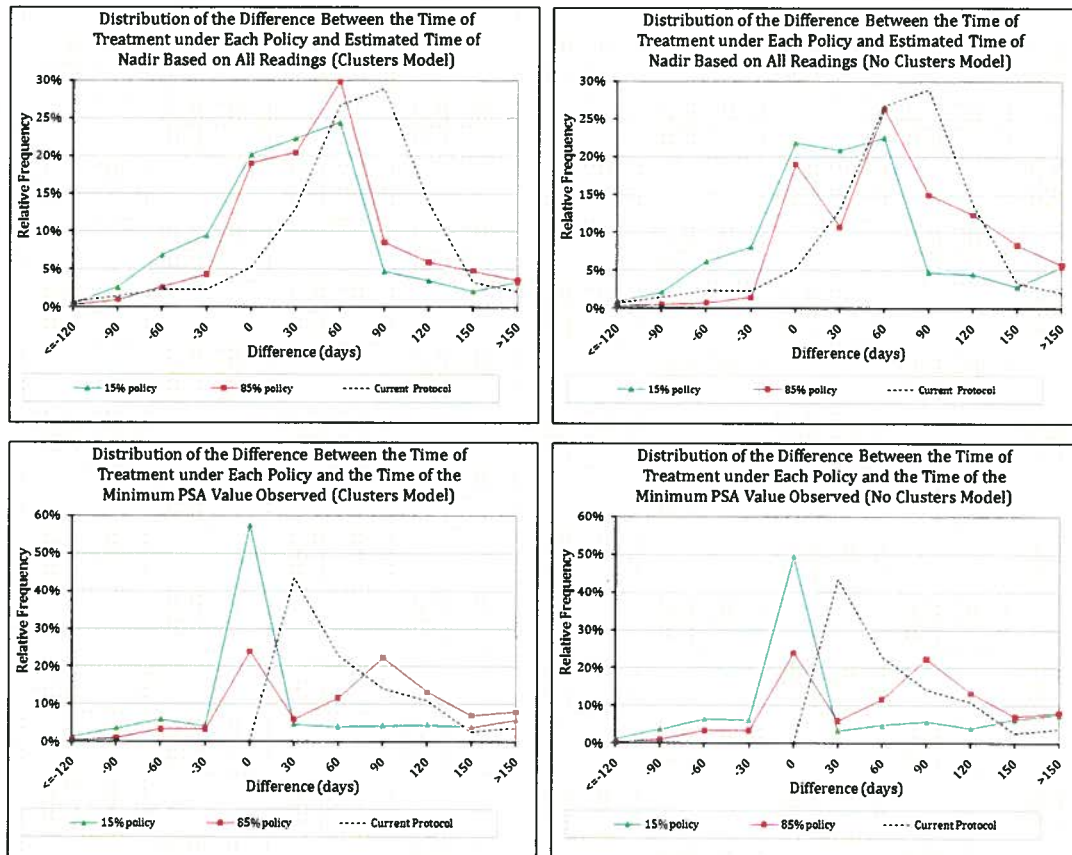
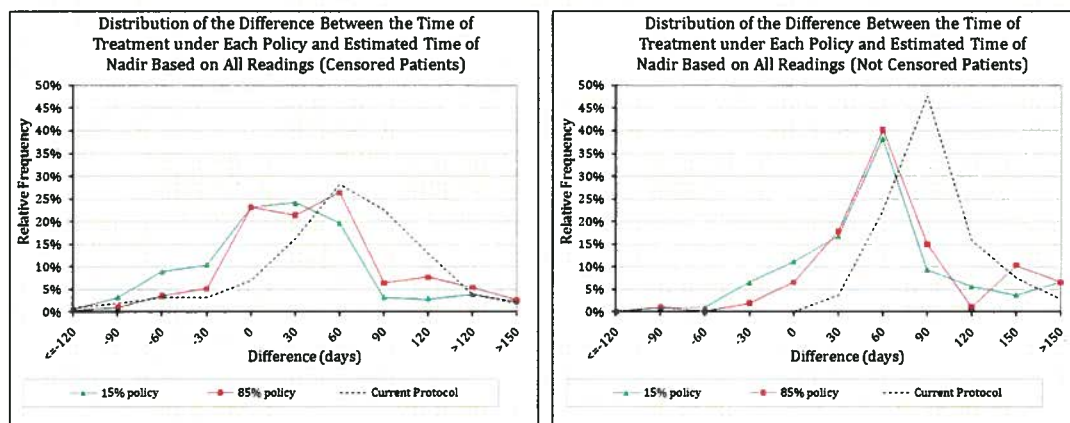


Figure 4.9 supports that, using our proposed policies, we are able to treat patients closer to the estimated nadir and to the time of the minimum PSA value observed than under the current protocol. Note that, by following the 15% policy, we could have treated over 75% of the patients within two months of the estimated nadir, while under the current protocol only 46% of the

patients received treatment within two months of the estimated nadir. Clustering has a bigger impact on the 85% policy than on the 15% policy.

Keeping the patient clusters, we now compare the distributions of the difference from the estimated time of nadir for censored and not censored patients (see Figure 4.10). Note that only the estimated time of nadir might be used for comparison in this case, since the time at which the minimum PSA value was observed for censored patients is highly dependent on when RT was started.

**Figure 4.10. Distribution of the difference from the estimated time of nadir for censored and not censored patients using the proposed policies and the current protocol**



For both censored and not censored patients the proposed policies outperform the current protocol. For not censored patients, the proportion of patients treated within two months of the estimated nadir is 73% and 66% using the 15% policy and the 85% policy, respectively, as compared to 26% under the current protocol. On the other hand, for censored patients this proportion is 77% and 76% as compared to 55% under the current protocol.

By recalibrating our clusters - as well as by investigating other thresholds - further improvements over the current protocol might be achieved. As highlighted in Chapter 3, the optimum patient specific policy might be obtained by formulating this problem as a Markov decision process. In addition, our policies assume that if the threshold probability is not reached by the end of eight months, all patients should commence RT. A modification of this policy might be to start RT if the threshold is reached or if the estimated time of nadir based on the updated parameters for a given patient is reached.



## 4.5 Conclusion and Remarks

In this chapter we broaden our analysis by applying the model developed in Chapter 3 to a larger, more heterogeneous cohort of patients. The model is used to decide when to start RT treatment of prostate cancer patients receiving neoadjuvant hormone therapy based on their PSA dynamics. To avoid bias, the model has been solely developed on a different set of patients to the ones used in our validation process and it is not recalibrated before testing. While further improvements should be achieved once the model has been recalibrated based on this new cohort of patients and an appropriate threshold for each patient is found, we show that, even with the imperfect model, we are able to predict the appropriate time of treatment for patients closer to when they reach their PSA nadir than under the current protocol.

Key to our analysis is the fact that, even at the time of RT treatment, it is not clear if and when the PSA nadir has been reached. We use two different estimates of the retrospective time of nadir: the time (or an interval around the time) at which the minimum PSA is observed and the estimated time of nadir obtained by fitting a curve through all observations. We favour the second estimate given the high proportion of censored patients - patients for whom the minimum PSA value observed occurs right before RT - yet both estimates are reported in this chapter. Regardless of the estimate used for comparison, our model outperforms the current protocol in terms of treating patients closer to the estimated nadir.

In addition, we show that, based on the prior information obtained from the other dataset and the initial PSA value for each patient, we are able to provide an estimate at the beginning of the hormone therapy of when patients are likely to reach their PSA nadir. Our results could be beneficial in forecasting demand and planning the (limited) RT resources. If at the beginning of the hormone treatment it is known that RT resources are not available to provide radiation when the patient is forecasted to reach his PSA nadir, clinicians may delay the start of the hormone treatment to ensure the patient is able to receive RT at his PSA nadir. The main tradeoffs of such decision must be investigated.

The accuracy of our model improves as new PSA readings become available, suggesting that taking PSA readings more frequently (especially at the beginning of the treatment) should be considered. Further research could study the optimal frequency of PSA readings.

We have clustered patients in our validation dataset based on the clustering algorithm developed using the initial dataset. While clustering patients does not significantly improve our nadir estimations, patients in the third cluster are more likely to require total androgen blockade until RT. By recalibrating our clustering algorithm based on this new cohort of patients, that likelihood might be further improved. Patients in the third cluster tend to have a higher initial PSA. Further analysis of the relationship between initial PSA and the type of hormone therapy required might be considered.

## CHAPTER 5: CONCLUSIONS, FUTURE CHALLENGES AND FINAL REMARKS

This thesis investigates and develops two healthcare operations research applications: health human resource planning and prostate cancer clinical decision making. Although both problems model dynamic decision making, they use different analytical approaches: optimization and large scale linear programming on one hand, and applied statistics, Kalman Filter and Bayesian updating on the other. Both models involve a multi-period planning horizon in which decisions have long term implications. I shall next summarize each contribution in more detail.

### 5.1 Main Conclusions

The workforce planning model determines the total number of students to admit to academic programs, the total number of nurses to train for management roles, and the total number of nurses and managers to recruit from outside the region to meet service needs. Our model incorporates consideration of “on-the-job” training and parental leaves (in the calculation of full time equivalence (FTE)), promotion rules, the existence of two types of education programs and the effect of an age specific attrition rate. To implement this model, a user friendly tool was developed to assist policy makers in planning the British Columbia registered nurses workforce. Its simplicity makes it ideal for scenario and “what-if” analyses. We demonstrate the capabilities of our model by investigating the impact of decreasing attrition from educational programs, changing nurse-to-manager ratios and exploring how other changes might alter planning recommendations.

Based on our analysis, we show that under the current nurse age distribution and attrition rates, only by having a large initial recruitment or by sacrificing the quality of service by changing nurse-to-population ratios, are we able to obtain a feasible solution to the problem. Practically, this implies that, unless proactive changes are made to the current workforce, the current level of service and nurse-to-manager ratios cannot be sustained in the future. As the current workforce is older than is optimal, adding RNs early on is necessary to meet future needs for RNs and managers. Given the importance of health human resources to ensure that healthcare needs of the population are met, changes to current practice need to be considered.

The model can be used to draw policy insights. Attrition from educational programs and from the profession have significant effects on recruitment and training. A key policy recommendation is that given the global shortage of nurses, as well as high recruitment costs, decision makers should consider initiatives that promote reductions of such attritions. Reducing attrition, especially by older nurses, for example by reducing nurse-manager ratios, and increasing the number of students admitted to advanced standing educational programs should also be considered. Nonetheless, decision makers should address constraints that limit the capacity of nursing educational programs to expand, such as the shortage of nursing faculty. On the other hand, a radical reduction in the duration of parental leaves would not have a significant effect on the solution. This may be a result of the low fertility rates in the province. However, if by increasing the length of parental leaves, attrition rates for nurses could be reduced, something that would have to be studied, this model would support such a decision. We emphasize that, to make a decision based on these results, decision makers need to consider other factors such as the possible association with the provincial fertility rate.

The patient-specific model challenges clinicians who treat patients with prostate cancer to think about patient's PSA nadirs and their predictions, rather than waiting for a rise in PSA or for eight months to have elapsed to start radiation therapy (RT) as is current practice. We are able to provide an estimate at the beginning of hormone therapy of when patients are likely to reach their PSA nadir. The accuracy of our predictions was validated by comparing our estimates with an independent sample's retrospective nadirs based on all PSA values and the time - or an interval around the time - at which the minimum PSA value for each patient was observed. Parameters used to model PSA progression are justified both theoretically and empirically. The distribution of the time of PSA nadir is derived from an approximation of the ratio of two correlated normal random variables. Estimates are updated as new patient information becomes available. Based on our analysis, we are able to predict the appropriate time of treatment closer to when patients reach their PSA nadir than is the case under current clinical practice. A user friendly tool is developed to assist clinicians in the decision to start RT based on the PSA progression of the patient.

Note that clustering patients did not significantly improve our nadir predictions. This might be due to the homogeneity of the patients used to develop the initial model. Moreover, patients in the third cluster had higher likelihood of requiring total androgen blockade until RT. Given the additional costs and co-morbidities associated with total androgen blockade, identifying which

patients are in this third cluster, and therefore more likely to require additional hormones during treatment, is of interest to clinicians. By recalibrating our clustering algorithm based on a more heterogeneous set of patients, cluster separation and predictions could be improved.

## 5.2 Future Challenges

Both models developed in this thesis lead to methodological challenges that could be further investigated.

### 5.2.1 The Workforce Planning Model

The workforce plan was affected by the fact that the model was only solved over a finite horizon. Future research should consider how to formulate appropriate terminal conditions in finite horizon approximations to infinite horizon models. While bounds were added to constrain the rate of change of number of admissions to educational programs in the last years, our solution nonetheless stipulates a decrease in admissions over this period. Various approaches have been discussed in the literature to address end effects in finite horizon approximations to infinite horizon models.

Lippman et al. (1967) and Kunreuther and Morton (1973) suggested that it might be possible to prove that decisions made in the first period of a multiple horizon problem will not be affected by values long enough into the future. The main advantage relies on the simplicity of this approach: it is only needed to solve the model over a horizon that is long enough to guide the decision(s) of the first period(s). However, the conditions for existence might be too restrictive, the length of the planning horizon might not be easy to find and, even if it exists, the planning horizon might be too long to be useful for the planner. See Puterman (1994) Section 6.8 for more on this approach in an MDP setting.

MacClain and Thomas (1977) in the context of production planning, on the other hand, suggested obtaining end conditions to the problem by solving a steady state problem in which no assumption is made with respect to the initial values of the decision variables, but an assumption is made to ensure that the model is in a steady state. The advantage of this approach is that, once the end condition is calculated, the problem may be solved over a shorter horizon than the planning horizon problem (good ending conditions might be better in summarizing future requirements

compared with adding more periods to the decision problem). The shorter period model might be easier for managers to understand. However, the authors assumed that it is always feasible to meet the end conditions, and it is unclear how the steady state conditions could be found in different settings; especially in our case when the model is driven by non-stationary input.

This second approach is favoured for the workforce planning model given the size of the original problem. As is shown in Chapter 2, the age distribution of the workforce had a significant impact on the model solution. It is hypothesized that placing bounds on the relative age distribution at the final period (either as constraints or as part of the objective function) might improve model results. The steady state model may be formulated by adding a constraint to avoid changes of the age distribution over time. Let  $FTE(n(1)_{ij})$ ,  $FTE(n(2)_{ij})$  and  $FTE(n(3)_{ij})$  be the total full time equivalent direct care nurses, entry level managers and senior level managers of age  $i$  in year  $j$ . Using the notation of Chapter 2, in the steady state model the age distribution constraints could be written as:

$$\frac{FTE(n(1)_{ij})}{\left(\frac{BCpop_j}{n(1)ratio_j}\right)} = \frac{FTE(n(1)_{ik})}{\left(\frac{BCpop_k}{n(1)ratio_k}\right)} \quad \forall i,j,k \quad (5.1)$$

$$\frac{FTE(n(2)_{ij})}{\left(\frac{BCpop_j}{n(1)ratio_j n(2)ratio_j}\right)} = \frac{FTE(n(2)_{ik})}{\left(\frac{BCpop_k}{n(1)ratio_k n(2)ratio_k}\right)} \quad \forall i,j,k \quad (5.2)$$

$$\frac{FTE(n(3)_{ij})}{\left(\frac{BCpop_j}{n(1)ratio_j n(2)ratio_j n(3)ratio_j}\right)} = \frac{FTE(n(3)_{ik})}{\left(\frac{BCpop_k}{n(1)ratio_k n(2)ratio_k n(3)ratio_k}\right)} \quad \forall i,j,k \quad (5.3)$$

Where the denominator in each formula corresponds to the minimum number of full time equivalent nurses needed in each period. We might look instead at the ratio of nurses between age categories, which should remain constant in the limit. Yet incorporating such ratios would involve nonlinear constraints and must be further investigated.

Boundary conditions to the original problem could be obtained from the age distribution of the steady state model. Future research should consider how to assign weights when the end conditions cannot be met by the original problem. In addition, if the steady state problem is to be solved over a finite horizon, it is challenging to prove that the steady state workforce plan is not affected by the planning horizon. An alternative might be to use a control theory approach in which the number of students admitted at time, the number of nurses recruited, and the number of nurses promoted are assumed to be our control variables.

### 5.2.2 The Prostate Cancer Model

In the case of the prostate cancer model - the question of how long to give neoadjuvant androgen ablation before starting radiotherapy is still open, as is the question of what is the best indicator for when to start radiotherapy. As discussed in Chapter 3, this problem may be formulated as a discrete time, finite horizon, Markov decision problem (MDP). Using the notation described in Chapter 3, the decision epochs correspond to the time  $t$  at which PSA readings are taken. The state of the model  $s$  corresponds to the patient's PSA curve parameters  $\hat{\theta}_{i,t|t-1}$  and the covariance of the parameters  $R_{i,t|t-1}$  used to calculate the distribution of the time of nadir. In each decision epoch, two possible actions may be taken: to start RT or to wait for another PSA reading. If RT is started, the patient receives a reward based on the probability of reaching the nadir in the following period, and no additional PSA readings are taken. The reward received can be interpreted as the probability of reaching the PSA nadir at the time RT is started. If the patient's clinician decides to wait, no reward is received. In that case, a new PSA reading is taken in the following period, and the process is repeated. The state is updated using the Kalman filter. Let  $t_{nadir}$  be the true time of nadir. If we assume that RT treatment is started within  $\xi$  days after the PSA reading is taken, and that the objective of the model is to maximize the probability that patients reach the PSA nadir in the period in which RT is started, the optimality equation  $v_t(\hat{\theta}_{i,t|t-1}, R_{i,t|t-1})$  might be written as:

$$v_t(\hat{\theta}_{i,t|t-1}, R_{i,t|t-1}) = \max \left\{ \begin{array}{l} P(t \leq t_{nadir} \leq t + \xi | \hat{\theta}_{i,t|t-1}, R_{i,t|t-1}), \\ \int P(\hat{\theta}_{i,t+1|t}, R_{i,t+1|t} | \hat{\theta}_{i,t|t-1}, R_{i,t|t-1}) v_t(\hat{\theta}_{i,t+1|t}, R_{i,t+1|t}) \end{array} \right\} \quad (5.4)$$

The objective function value can be interpreted as the maximum probability of reaching the PSA nadir when RT is started given the curve parameters and expected PSA readings and parameter updates. This is an optimal stopping problem. The first term in (5.4) represents the distribution of the time of nadir, while the second term is the usual dynamic recursion which can be obtained from the Kalman Filter results. Other alternatives to the objective, such as maximizing the patient's quality of life, might be considered. Moreover, clusters of relatively homogenous sets of patients might be incorporated in this decision model.

This model entails a continuous, partially observable state space model, which is difficult to solve in practice. However, as discussed in Chapter 3, the parameter covariances do not depend on the new PSA readings and can be computed off-line based on the prior information. In addition, while

the reward function is based on the ratio of correlated normal random variables, the integral approximation discussed would simplify the reward calculations. The state space might be discretized to obtain initial solutions to the model, and by analyzing properties of the reward function, policy implications might be obtained.

To implement the patient specific model, our first step has been to test it on other databases to validate the PSA curve dynamics and nadir predictions, accuracy improvements over time, cluster classification and the policy implications of our model. The next step involves testing the tool by conducting a pilot study in which oncologists are asked to provide qualitative feedback on the tool and to suggest possible improvements to facilitate its use. A formal randomized clinical trial should be formulated to validate the decision model. The purpose of the trial would be to compare the current protocol to the proposed model in terms of survival and time to PSA relapse. A cohort of intermediate/high risk patients would receive treatment under the current protocol, while another cohort of intermediate/high risk patients would receive treatment following the best policies obtained from our validation database. Patients should be monitored at least once a month during hormone treatment and at least once every six months after radiation therapy is given. Comparisons could then be made at the end of three, five and ten years.

While decisions in the workforce planning model are made regularly (once a year), decisions at the disease management level are made as new information becomes available, that is, after each PSA reading is taken. Future research should consider how to incorporate the decision of when to take additional PSA readings based on observed PSA progression in the MDP framework. This may be formulated as a discrete time MDP in which the state of the model correspond to the patient's PSA curve parameters and the covariance of the parameters, yet the possible actions are whether to take an additional reading or not in each period and whether to start RT based on the knowledge of the PSA dynamics that the clinician is expected to have at the end of the period. Each period corresponds to the possible times readings could be taken. Only if an additional PSA reading is taken in that period is the curve updated. A specific challenge would be developing an appropriate representation for costs so that model objectives are on the same scale.

### **5.3 Concluding Remarks**

While in the workforce planning model decisions are made at the beginning of the planning horizon, in the prostate cancer model decisions incorporate additional information obtained as



new data (PSA readings) become available. Limited availability and accessibility of data played a key role in the development of both models. Greater emphasis should be placed on accurate collection and calculation of the parameters that have the greatest effect on the decision-making process. Our workforce planning analysis points to the need to obtain better estimations of attrition rates, both from educational programs and from the profession. The prostate cancer treatment model suggests that additional PSA readings may have produce significant improvements in the nadir estimations. Taking additional readings, especially in the first months of hormone treatment, might not only improve nadir estimation, but may be beneficial in forecasting demand and thus improve planning for the use of limited RT resources.

By modeling healthcare problems faced at both a strategic level and a patient specific level, I have gained a deeper understanding of the benefits and challenges of applying operations research in the field of healthcare. The next step is to link strategic decision making with patient-specific models. Policy decisions might influence how patients are treated, and treatment decisions will influence the resources required. For instance, the level of resources might impact the  $\xi$  used in (5.4), and the number of patients needing treatment will impact the necessary resources in constraints (2.23), (2.24) and (2.25). Future research shall combine both types of models in order to compare optimum policies attained from a patient's perspective to system-wide policies that take resource availability into consideration by estimating expected deviations from the optimum time of treatment based on the levels of resources available.

To conclude, I believe healthcare operations research to be a very promising area of research. Operations research can be used to guide decision makers in making better policy, operations management and disease management decisions. Those decisions are not only significant from the patient and economic point of view, but they involve challenging methodological questions that need to be addressed. My hope is that the field of healthcare operations research will continue to grow, as more operational researchers and healthcare professionals acknowledge the significance and potential of such collaboration.

## REFERENCES

- Abrahams, C. 2005. *In Search of the 'Right' Ontario Postgraduate Allocation: The ADIN Model for HHR Planning*. Ontario: Ministry of Health and Long-Term Care.
- Agresti, A. 1996. "An Introduction to Categorical Data Analysis" New York, NY: John Wiley & Sons, Inc.
- Alagoz, O., L. Maillart, A. Schaefer and M. Roberts. 2000. "The Optimal Timing of Living-Donor Liver Transplantation." *Management Science*. 46(9): 1420-1430.
- American Cancer Society. 2008. "Detailed Guide: Prostate Cancer" Atlanta, Georgia: American Cancer Society, Retrieved June 30, 2009  
([http://www.cancer.org/docroot/CRI/CRI\\_2\\_3x.asp?dt=36](http://www.cancer.org/docroot/CRI/CRI_2_3x.asp?dt=36)).
- American Joint Committee on Cancer. 2002. *Cancer Staging Handbook. From the AJCC Cancer Staging Manual*. Sixth edition, edited by F.L. Greene, D.L. Page, I.D. Fleming, A.G. Fritz, C.M. Balch, D.G. Haller and M. Morrow. Springer. New York.
- Archambault, R. 1999. *New COPS Occupational Projection Methodology*. Human Resources Development Canada. Retrieved August 4, 2006  
([www11.hrsdc.gc.ca/en/cs/sp/hrsdcrarb/publications/research/1999-000135/page00.shtml](http://www11.hrsdc.gc.ca/en/cs/sp/hrsdcrarb/publications/research/1999-000135/page00.shtml)).
- Aykin, T. 1996. "Optimal Shift Scheduling with Multiple Break Windows" *Management Science*. 42(4): 591-602.
- B.C. Stats, B.C. 2003 (June 6). "Occupational Employment Projections 2001-2011." *Labour Force Statistics*.
- Baker, R. 1998. "Use of a Mathematical Model to Evaluate Breast Cancer Screening Policy." *Health Care Management Science*. 1(2): 103-113.
- Balinsky W. and A. Reisman. 1972. "Some Manpower Planning Models Based on Levels of Educational Attainment." *Management Science*. 18(12): 691-705.

- Basu, K. and J. Liu. 2005. "Registered Nurse Demand Model." *Health Human Resources' Modelling Working Group Workshop*. Microsimulation Modelling and Data Analysis Division Health Canada.
- Basu, K. and K. Halliwell. 2004. "Projecting the HHR Impacts of Demographic Change." *Health Policy Research Bulletin*. May 4, 2004: 17-21.
- Baumann, A., L. O'Brien-Pallas, M. Armstrong-Stassen, J. Blythe, R. Bourbonnais, S. Cameron, D. Irvine Doran, M. Kerr, L. McGillis Hall, M. Vezina, M. Butt, and L. Ryan. 2001. *Commitment and Care: The Benefits of a Healthy Workplace for Nurses, their Patients and the System. A Policy Synthesis*. Ottawa: Canadian Health Services Research Foundation.
- BC Cancer Agency. 2009. "Prostate Cancer: BC Cancer Agency" Vancouver, BC: BC Cancer Agency, Retrieved June 30, 2009  
(<http://www.bccancer.bc.ca/>).
- BC Stats. 2006. "Projections - Total Age Group - Query Results" *Population Projections - BC and Regional*. BC Stats. Retrieved November 28, 2006  
([www.bcstats.gov.bc.ca/data/pop/pop/popproj.asp#admin](http://www.bcstats.gov.bc.ca/data/pop/pop/popproj.asp#admin)).
- Brandeau, M. L., F. Sainfort and W. P. Pierskalla. 2004. "Health Care Delivery: Current Problems and Future Challenges" Chapter 1 of M. L. Brandeau et. al, *Operations Research and Health Care. A Handbook of Methods and Applications*, Kluwer Academic Publishers.
- Butler, D. 1979. "A Hazardous-Inspection Model." *Management Science*. 25(1): 79-89.
- Canadian Cancer Society's Steering Committee. 2009. "Canadian Cancer Statistics 2009" Toronto, Ontario: Canadian Cancer Society, Retrieved June 30, 2009  
(<http://www.cancer.ca/statistics>).
- Canadian Health Services Research Foundation. 2004. "Listening for Direction: A National Consultation on Health Services and Policy Issues" Ottawa, Ontario: Canadian Health Services Research Foundation. N.d.  
([http://www.chsrf.ca/other\\_documents/listening/](http://www.chsrf.ca/other_documents/listening/)).
- Canadian Institute for Health Information (CIHI). 2006. *Workforce Trends of Registered Nurses in Canada, 2005*. Ottawa: CIHI.

- Canadian Institute for Health Information (CIHI). 2007. *Canada's Health Care Providers, 2007*. Ottawa: CIHI.
- Canadian Institute for Health Information (CIHI). 2008. *Health Care in Canada 2008*. Ottawa, Ontario: CIHI.
- Canadian Nurses Association. 2002. *Planning for the Future: Nursing Human Resource Projections*. Ottawa: Canadian Nurses Association.
- Celeux, G., and G. Soromenho. 1996. "An Entropy Criterion for Assessing the Number of Clusters in a Mixture Model" *Journal of Classification* 13: 195-212.
- Chattopadhyay, A.K. and A. Gupta. 2007. "A Stochastic Manpower Planning Model under Varying Class Sizes." *Annals of Operations Research*. 155: 41-49.
- Chick, S., S. Soorapanth, and J. Koopman. 2003. "Inferring Infection Transmission Parameters That Influence Water Treatment Decisions." *Management Science*. 49(7): 920-935.
- College of Registered Nurses of British Columbia. 2006. *New Graduate Registered Nurse Study – 2005*. Vancouver: College of Registered Nurses of British Columbia.
- D'Amato, R., R. D'Aquila, L. Wein. 2000. "Management of Antiretroviral Therapy for HIV Infection: Analyzing When to Change Therapy" *Management Science* 50 (10): 1200-1213.
- Daskin, M. S. and L. Dean. 2004. "Location of Health Care Facilities" Pp. 43-76 in *Operations Research and Health Care. A Handbook of Methods and Applications*, edited by M. L. Brandeau, F. Sainfort, and W. P. Pierskalla. Norwell, MA : Kluwer Academic Publisher.
- Dault, M., J. Lomas, and M. Barer on behalf of the Listening for Direction II partners. 2004. Listening for Direction II: National Consultation on Health Services and Policy Issues for 2004-2007: Final Report. Ottawa: Canadian Health Services Research Foundation. Retrieved June 25, 2007  
([http://www.chsrf.ca/other\\_documents/listening/index2\\_e.php/](http://www.chsrf.ca/other_documents/listening/index2_e.php/)).
- De Feyter, T. 2007. "Modeling Mixed Push and Pull Promotion Flows in Manpower Planning." *Annals of Operations Research*. 155(1): 25-39.

- Dickson, G. L. 2002. "Forecasting the Demand for Nurses in New Jersey." New Jersey Collaborating Center for Nursing. N.d.  
([www.njccn.org/pdf/forecast\\_NJ.pdf](http://www.njccn.org/pdf/forecast_NJ.pdf)).
- Felici, G. and C. Gentile. 2004. "A Polyhedral Approach for the Staff Rostering Problem" *Management Science*. 50(3): 381-393.
- Ferris, M., J. Lim, and D. Shepard. 2003. "An Optimization Approach for Radiosurgery Treatment Planning." *SIAM Journal on Optimization*. 13(3): 921-937.
- Fieller, E. C. 1932. "The distribution of the index in a normal bivariate population" *Biometrika* 24, 428-440.
- Fisher, M., K. Ramdas, and Y. Zheng. 2001. "Ending Inventory Valuation in Multiperiod Production Scheduling." *Management Science*. 47(5): 679-692.
- Fraley, C., and A. E. Raftery. 2002. "Model-Based Clustering, Discriminant Analysis, and Density Estimation" *Journal of the American Statistical Association* 97: 611-631.
- Fraley, C., and A. E. Raftery. 1998. "How Many Clusters? Which Clustering Method? Answers Via Model-Based Cluster Analysis" *The Computer Journal* 41 (8): 578-588.
- Fraley, C., and A. E. Raftery. 2006. "MCLUST Version 3 for R: Normal Mixture Modeling and Model-based Clustering" *Technical Report* No. 504. Seattle, Washington: Department of Statistics, University of Washington.
- Frontline Systems, Inc. 2007. *Optimization Software for Excel, .NET, Java, MATLAB*. Incline Village, USA. N.d.  
([www.solver.com](http://www.solver.com)).
- Fuloria, P. and S. Zenios. 2001. "Outcomes-Adjusted Reimbursement in a Health-Care Delivery System." *Management Science*. 47(6): 735-751.
- Georgiou, A.C. and N. Tsantas. 2002. "Modelling Recruitment Training in Mathematical Human Resource Planning." *Applied Stochastic Models in Business and Industry*. 18: 53-74.

- Gleave, M. E., S. La Bianca, and S. L. Goldenberg. 2000. "Neoadjuvant Hormonal Therapy Prior to Radical Prostatectomy: Promises and Pitfalls" *Prostate Cancer and Prostatic Diseases* 3: 136-144.
- Goldie, J. H., and A. J. Coldman. 1979. "A Mathematic Model for Relating the Drug Sensitivity of Tumors to their Spontaneous Mutation Rate" *Cancer Treatment Reports* 63: 11-12, 1727-1733.
- Green, L. 2004. *Capacity Planning and Management in Hospitals, in: Operations Research and Health Care. A Handbook of Methods and Applications*, edited by M.L. Brandeau, F. Sainfort and W. Pierskalla. Kluwer Academic Publishers.
- Grunow, M., H. Günther and G. Yang. 2004. "Development of a Decision Support Model for Scheduling Clinical Studies and Assigning Medical Personnel" *Health Care Management Science*. 7(4): 305-317.
- Güne, E., S. Chick, and O. Akin. 2004. "Breast Cancer Screening Services: Trade-offs in Quality, Capacity, Outreach, and Centralization." *Health Care Management Science* 7(4): 291-303.
- Hall, N., J. Hershey, L. Kessler, and R. Stotts. 1992. "A Model for Making Project Funding Decisions at the National Cancer Institute." *Operations Research*. 40(6): 1040-1052.
- Harper, P. and S. Jones. 2005. "Mathematical Models for the Early Detection and Treatment of Colorectal Cancer." *Health Care Management Science*. 8(2): 101-109.
- Harvey, A. 1991. "State space models and the Kalman filter" Pp. 100-167 in *Forecasting Structural Time Series Models and the Kalman Filter*. Cambridge, UK: Cambridge University Press.
- Henderson, S. and A. Mason. 2004. "Ambulance Service Planning: Simulation and Data Visualisation." Pp. 77-102 in *Operations Research and Health Care. A Handbook of Methods and Applications*, edited by M. L. Brandeau, F. Sainfort, and W. P. Pierskalla. Norwell, MA : Kluwer Academic Publisher.
- Heymann, J. J., M. C. Benson, K. M. O'Toole, B. Malyszko, R. Brody, D. Vecchio, P. B. Schiff, M. M. Mansukhani, and R. D. Ennis. 2007. "Phase II Study of Neoadjuvant Androgen Deprivation Followed by External-Beam Radiotherapy With 9 Months of Androgen Deprivation for

- Intermediate- to High-Risk Localized Prostate Cancer" *Journal of Clinical Oncology* 25(1): 77-84.
- Horwitz, E. M., and G. E. Hanks. 2000. "Hormonal Therapy Combined with Radiation or Surgery in Treatment of Locally Advanced Non-metastatic Prostate Cancer" Pp. 782-789 in *Comprehensive Textbook of Genitourinary Oncology*, edited by N. J. Vogelzang, P. T. Scardino, W. U. Shipley, D. S. Coffey and B. J. Miles. Philadelphia, PA: Lippincott Williams & Wilkins.
- Hu, C., W. Lovejoy and S. Shafer. 1996. "Comparison of Some Suboptimal Control Policies in Medical Drug Therapy." *Operations Research*. 44(5): 696-709.
- Human Resources and Skills Development Canada (HRSDC) B.C./Yukon Region and B.C. Ministry of Advanced Education. 2005. "Work Futures B.C. Occupational Outlooks: Registered Nurses (NOC 3152)." Retrieved September 25, 2007 (<http://www.workfutures.bc.ca/profiles/profile.cfm?noc=3152&lang=en&site=graphic>).
- Iravani, S.M.R., B. Kolfal, and M. P. Van Oyen. 2007. "Call-Center Labor Cross Training: It's a Small World After All." *Management Science*. 53(7): 1102 - 1112.
- Kazanjian, A., I.R. Pulcins, and K. Kerluke. 1992. "A Human Resources Decision Support Model: Nurse Deployment Patterns in One Canadian System." *Hospital and Health Services Administration*. 37(3): 303-319.
- Kazanjian, A., K. Brothers, and G. Wong. 1986. "Modeling the Supply of Nurse Labor. Life-cycle Activity Patterns of Registered Nurses in one Canadian Delivery System." *Medical Care*. 24(12): 1067-1083.
- Keehan, S., A. Sisko, C. Truffer, S. Smith, C. Cowan, J. Poisal, M. K. Clemens, National Health Expenditure Accounts Projections Team. 2008. "Health Spending Projections Through 2017." *Health Affairs - Web Exclusive*. 27(2): w145-w155.
- Kephart, G., S. Maaten, L. O'Brien-Pallas, G. Tomblin Murphy and B. Milburn. 2004. "Simulation Analysis Report" *Building the Future: An Integrated Strategy for Nursing Human Resources in Canada*. Ottawa. Retrieved August 4, 2006 ([www.buildingthefuture.ca/e/study/phase1/reports/](http://www.buildingthefuture.ca/e/study/phase1/reports/)).

- Kirch, R. and M. Klein. 1974. "Surveillance Schedules for Medical Examinations." *Management Science*. 20(10): 1403-1409.
- Knobbe, E. J., and B. Buckingham. 2005. "The extended Kalman filter for continuous glucose monitoring" *Diabetes Technology & Therapeutics* 7(1): 15-27.
- Kognakorn, T. and F. Sainfort. 2004. "Modeling Health Outcomes for Economic Analysis." Pp. 255-274 in *Operations Research and Health Care. A Handbook of Methods and Applications*, edited by M. L. Brandeau, F. Sainfort, and W. P. Pierskalla. Norwell, MA : Kluwer Academic Publisher.
- Kramer, M. and C. Schmalenberg. 2002. "Staff Nurses Identify Essentials of Magnetism." *Magnet Hospitals Revisited: Attraction and Retention of Professional Nurses*, edited by M. McClure and A. Hinshaw. Washington, DC: American Nurses Publishing.
- Kramer, M., C. Schmalenberg, and P. Maguire. 2004. "Essentials of a Magnetic Work Environment. Part 3." *Nursing*. 34(8): 44-47.
- Kunreuther H.C., and T.E. Morton. 1973. "Planning Horizons for Production Smoothing with Deterministic Demands." *Management Science*. 20 (1): 110-125.
- Laporte, G., F. Semet, V. Dadeshidze, and L. Olsson. 1998. "A Tiling and Routing Heuristic for the Screening of Cytological Samples." *The Journal of the Operational Research Society*. 49(12): 1233-1238.
- Lincoln, T. and G. Weiss. 1964. "A Statistical Evaluation of Recurrent Medical Examinations." *Operations Research*. 12(2): 187-205.
- Lippman, S. A., A. J. Rolfe, H. M. Wagner, and J. S. C. Yuan. 1967. "Algorithms for Optimal Production Scheduling and Employment Smoothing" *Operations Research*. 15 (6): 1011-1039.
- Lovejoy, W. and Y. Li. 2002. "Hospital Operating Room Capacity Expansion." *Management Science*. 48(11): 1369-1387.
- Makarov, D. V., B. J. Trock, E. B. Humphreys, L. A. Mangold, P. C. Walsh, J. I. Epstein and A. W. Partin. 2007. "Updated Nomogram to Predict Pathologic Stage of Prostate Cancer Given Prostate-specific Antigen Level, Clinical Stage, and Biopsy Gleason Score (Partin Tables) Based on Cases from 2000 to 2005." *Urology*. 69(6): 1095-1101.



- Mangasarian, O., W. Street and W. Wolberg. 1995. "Breast Cancer Diagnosis and Prognosis via Linear Programming." *Operations Research*. 43(4): 570-577.
- Martel, A. and W. Price. 1981. "Stochastic Programming Applied to Human Resource Planning." *The Journal of the Operational Research Society*. 32(3): 187-196.
- McClain, J. and J. Thomas. 1977. "Horizon Effects in Aggregate Production Planning with Seasonal Demand." *Management Science*. 23(7): 728-736.
- Mehrez, A. and A. Gafni. 1987. "The Optimal Treatment Strategy - A Patient's Perspective." *Management Science*. 33(12): 1602-1612.
- Morris, C. N. 1983. "Parametric Empirical Bayes Inference: Theory and Applications", *Journal of the American Statistical Association* 78(381): 47-55.
- Morris, W. J. et al. 2009. "Androgen Suppression Combined With Elective Nodal and Dose Escalated Radiation Therapy" NCI Clinical Trials Inventory, Retrieved April 30, 2009  
(<http://www.cancer.gov/search/ViewClinicalTrials.aspx?cdrid=453036&version=HealthProfessional&protocolsearchid=6093623>).
- Newton, S. and L. Buske. 1998. "Physician Resource Evaluation Template: A Model for Estimating Future Supply in Canada" *Annals of the Royal College of Physicians and Surgeons of Canada*. 31(3): 145-50.
- Niblock, P., T. Pickles, and the British Columbia Cancer Agency Prostate Cohort Outcomes Initiative. 2006. "Rising Prostate-Specific Antigen Values During Neoadjuvant Androgen Deprivation Therapy: the Importance of Monitoring" *International Journal of Radiation Oncology Biology Physics* 65(1): 59-64.
- O'Brien-Pallas, L., A. Baumann, G. Donner, G. Tomblin Murphy, J. Lochhaas-Gerlach, and M. Luba. 2001. "Forecasting Models for Human Resources in Health Care" *Journal of Advanced Nursing*. 33(1): 120-29.
- O'Brien-Pallas, L., C. Alksnis, and S. Wang. 2003. *Bringing the Future into Focus. Projecting RN Retirement in Canada*. Ottawa: CIHI.

- Organizing for America. 2009. "Healthcare. The Current Situation" Washington, DC: Organizing for America, Retrieved July 27, 2009 (<http://www.barackobama.com>).
- Özekici, S. and S. Pliska. 1991. "Optimal Scheduling of Inspections: A Delayed Markov Model with False Positives and Negatives." *Operations Research*. 39(2): 261-273.
- Puterman, M. 1994. "Convergence of Policies, Turnpikes, and Planning Horizons" Pp. 218-223 in *Markov Decision Processes*. New York, New York: John Wiley and Sons.
- Pardalos, P., W. Chaovalitwongse, L. D. Iasemidis, J. C. Sackellares, D. Shiau, P. R. Carney, O. A. Prokopyev and V. A. Yatsenko. 2004. "Seizure Warning Algorithm Based on Optimization and Nonlinear Dynamics", *Mathematical Programming*. 101(2): 365-385.
- Parker, R.S., F.J. Doyle III, and N.A. Peppas. 1999. "A Model-Based Algorithm for Blood Glucose Control in Type I Diabetic Patients" *Biomedical Engineering, IEEE Transactions* 46(2): 148-157.
- Partin, A. W., and B. Carter. 2000. "Use of Prostate-Specific Antigen Density and Velocity" Pp. 650-654 in *Comprehensive Textbook of Genitourinary Oncology*, edited by N. J. Vogelzang, P. T. Scardino, W. U. Shipley, D. S. Coffey and B. J. Miles. Philadelphia, PA: Lippincott Williams & Wilkins.
- Patrick, W.J. 2001. "Estimating First-Year Student Attrition Rates: An Application of Multilevel Modeling Using Categorical Variables." *Research in Higher Education*. 42(2): 151-170.
- Peterson, M. 1996. "Two Models for Assessing a Federal Environmental Health Policy: The Case of Radon in U.S. Homes." *Management Science*. 42(10): 1476-1492.
- Pringle, D. and L. Green. 2005. "The Pulse of Renewal: A Focus on Nursing Human Resources. Examining the Causes of Attrition from Schools of Nursing in Canada" *Canadian Journal of Nursing Leadership*. May 2005 (<http://www.longwoods.com/view.php?id=&aid=17477&cat=130>).

- Pringle, D.M., L. Green and S. Johnson. 2004 (December). "Nursing Education in Canada: Historical Review and Current Capacity." In L. O'Brien-Pallas, A. Baumann, L. Piazza, D. Pringle, G. Tomblin Murphy, S. Birch, G. Kephart, L. Nagle, L. McGillis Hall, G. Donner, G. Pink, D. Thomson, M. Gunderson, M. Lemonde, C. Alksnis, I. Zeytinolu and J. Blythe, eds., *Building the Future: An Integrated Strategy for Nursing Human Resources in Canada*. Ottawa: The Nursing Sector Study Corporation. Retrieved December 5, 2006 (<http://www.buildingthefuture.ca/e/study/phase1/reports/Step8-English-Final-July8.pdf>).
- Rao, P. 1990. "A Dynamic Programming Approach to Determine Optimal Manpower Recruitment Policies." *The Journal of the Operational Research Society*. 41(10): 983-988.
- Rehbinder, H., C. Forsgren, and J. Löf. 2004. "Adaptive Radiation Therapy for Compensation of Errors in Patient Setup and Treatment Delivery" *Medical Physics* 31(12): 3363-3371.
- Romanow, R.J. 2002. *Building on Values: The Future of Health Care in Canada: Final Report*. Commission on the Future of Health Care in Canada.
- Romeijn, H., R. Ahuja, J. Dempsey, and A. Kumar. 2005. "A Column Generation Approach to Radiation Therapy Treatment Planning Using Aperture Modulation." *SIAM Journal on Optimization*. 15(3): 838-862.
- Sanchez de Rivera, D. 1980. "A Decision Analysis Model for a Serious Medical Problem." *Management Science*. 26(7): 707-718.
- Sandblom, G., L. Holmberg, J. E. Damberg, J. Hugosson, J. E. Johansson, R. Lundgren, E. Mattsson, J. Nilsson and E. Varenhorst. 2002. "Prostate-Specific antigen as Surrogate for Characterizing Prostate Cancer Subgroups" *Scandinavian Journal of Urology & Nephrology*. 36(2): 106-112.
- Santibañez, P., M. Begen, and D. Atkins. 2007. "Surgical Block Scheduling in a System of Hospitals: an Application to Resource and Wait List Management in a British Columbia Health Authority." *Health Care Management Science*. 10(3): 269-282.
- Scott, G. 2004. "Fines for High Student Attrition Rates." *Nursing Standard*. 19(9): 5.
- Shechter, S. M., M. D. Bailey, A. J. Schaefer and M.S. Roberts. 2008. "The Optimal Time to Initiate HIV Therapy Under Ordered Health States." *Operations Research*. 56(1): 20-33.

- Sherlaw-Johnson, C. and S. Gallivan. 2000. "The Planning of Cervical Cancer Screening Programmes in Eastern Europe: is Viral Testing a Suitable Alternative to Smear Testing?" *Health Care Management Science*. 3(4): 323-329.
- Shwartz, M. 1978. "A Mathematical Model Used to Analyze Breast Cancer Screening Strategies." *Operations Research*. 26(6): 937-955.
- So, K. and C. Tang. 2000. "Modeling the Impact of an Outcome-Oriented Reimbursement Policy on Clinic, Patients, and Pharmaceutical Firms." *Management Science*. 46(7): 875-892.
- Statistics Canada. 2004. Provinces and Territories. "Live Birth Rate, Age-Specific, per 1000 Females in Age Range." Ottawa: Statistics Canada.
- Task Force Two. 2005. *A Physician Human Resource Strategy for Canada: Innovation Service Models in Canada Database*. Ottawa: Government of Canada.
- Thompson, P. 2007. "How Much Did the Liberty Shipbuilders Forget?" *Management Science*. 53(6): 908 - 918.
- Tomblin Murphy, G. 2005. "Health Human Resource Planning: An Examination of Relationships Among Nursing Service Utilization, and Estimate of Population Health, and Overall Health Outcomes in the Province of Ontario." Toronto, Ontario: University of Toronto. PhD Thesis.
- Tomblin Murphy, G. and O'Brien-Pallas, L. 2004. "How do Health Human Resources Policies and Practices Inhibit Change? A Plan for the Future." *Changing Health Care in Canada, The Romanow Papers, Volume 2*, edited by P.G. Forest, G.P. Marchildon, and T. McIntosh. Toronto, Ontario: University of Toronto Press.
- Vanhoucke, M. and B. Maenhout. 2007. "NSPLIB – A Nurse Scheduling Problem Library: a Tool to Evaluate (meta-) Heuristic Procedures." *Operational Research for Health Policy: Making Better Decisions*, edited by S. Brailsford and P. Harper. Bern: Peter Lang. Proceedings of the 31st Annual Conference of the European Working Group on Operational Research Applied to Health Services.
- Verter, V. and S. Lapierre. 2002. "Location of Preventive Health Care Facilities." *Annals of Operations Research*. 110(1-4): 123-132.

- Villeneuve, M. and J. MacDonald. 2006. *Towards 2020: Visions for Nursing*. Ottawa: Canadian Nurses Association.
- Weber, N. 2005. *The Hidden Costs of RN Turnover*. New York State Nurses Association. Retrieved September 25, 2008  
(<http://www.nysna.org/publications/report/2005/feb/turnover.htm>).
- Wharrad, H.J., M. Chapple, and N. Price. 2003. "Predictors of Academic Success in a Bachelor of Nursing Course." *Nurse Education Today*. 23(4): 246-54.
- Wiersma, E. 2007. "Conditions That Shape the Learning Curve: Factors That Increase the Ability and Opportunity to Learn" *Management Science*. Retrieved online before print Oct 19, 2007  
(<http://mansci.journal.informs.org/cgi/content/abstract/53/12/1903>).
- Zhu, H., and R. Rohwer. 1996. "Bayesian regression filters and the issue of priors", *Neural Computing & Applications*. 4(3): 130-142.