

# WHEN TO TREAT PROSTATE CANCER PATIENTS BASED ON THEIR PSA DYNAMICS

**MARIEL S. LAVIERI**

*Industrial and Operations Engineering, University of Michigan, 1205 Beal Avenue, Ann Arbor, MI, 48109, USA [lavieri@umich.edu](mailto:lavieri@umich.edu)*

**MARTIN L. PUTERMAN**

*Sauder School of Business, The University of British Columbia, 2053 Main Mall, Vancouver, BC, V6T 1Z2, Canada [marty@chcm.ubc.ca](mailto:marty@chcm.ubc.ca)*

**SCOTT TYLDESLEY, MD**

*Radiation Oncology, British Columbia Cancer Agency, 600 West 10th Avenue, Vancouver, BC, V5Z 4E6, Canada [styldesl@bccancer.bc.ca](mailto:styldesl@bccancer.bc.ca)*

**WILLIAM J. MORRIS, MD**

*Radiation Oncology, British Columbia Cancer Agency, 600 West 10th Avenue, Vancouver, BC, V5Z 4E6, Canada [jmorris@bccancer.bc.ca](mailto:jmorris@bccancer.bc.ca)*

We present a novel approach to model individual disease progression of prostate cancer patients who receive hormone therapy before radiation therapy. Our model is used to decide when to initiate radiation therapy based on the 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 to 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. Our model trades off the risk of beginning radiation therapy too soon so that hormone therapy has not achieved its maximum effect vs. waiting too long to start therapy so that there is an increased risk of tumor cells becoming resistant to the treatment. We illustrate 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 in each patient.

## 1. Introduction

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 on our study, resulting in a short data series for each patient. By the time sufficient patient specific data is 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 is gathered sequentially, so as to make better decisions regarding the best time to initiate RT.

According to the National Cancer Institute (2009), prostate cancer has the highest incidence and is the second leading cause of cancer death in men. In 2009, an estimated 192,280 men will be diagnosed with prostate cancer 27,360 will die of the disease in the United States. 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.

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, La Bianca and Goldenberg 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 the patient’s 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. 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 and/or that some of the cells have become resistant to the hormone treatment (Gleave, La Bianca and Goldenberg 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).

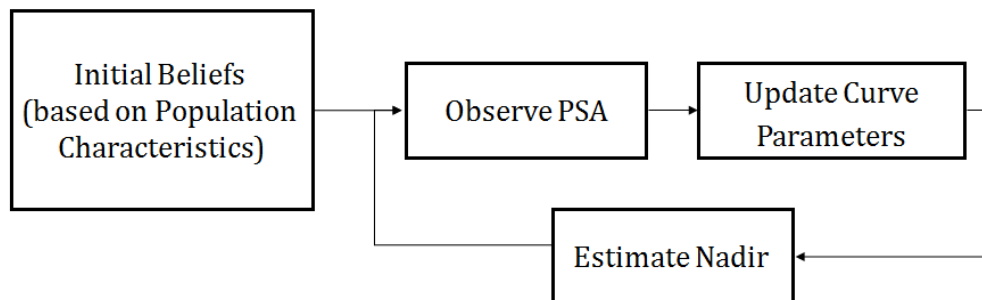
Gleave, La Bianca and Goldenberg (2000) hypothesized that “maximal tumor regression probably occurs when PSA reaches its nadir level”. We assume that it is ideal to start RT when the PSA reaches its minimum level. However, since cancer progresses and regresses at different rates in patients, it is difficult to predict when such 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. Alagoz et al. (2007) modeled the decision of when to accept a cadaveric liver for transplantation. Shechter et al. (2008) addressed the question of when to initiate HIV treatment. Both models rely on use of discrete health states, which might be difficult to define in practice. We do not use discrete health states in our model. Instead, we model probabilities and their dynamics. The quality of the decision is expressed in terms of the probability of treating at the “right time” (when PSA reaches its minimum level).

D’Amato, D’Aquila and Wein (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. 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. 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 1.

**Figure 1. Model overview.**



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.

The paper is organized as follows. The PSA dynamics model is described in Section 2.1. Section 2.2 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 2.3. Section 2.4 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 2.5. We illustrate our modeling approach in Section 3 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 4. A summary of the notation used appears in Appendix A.

## 2. Methods

### 2.1 Modeling PSA Dynamics

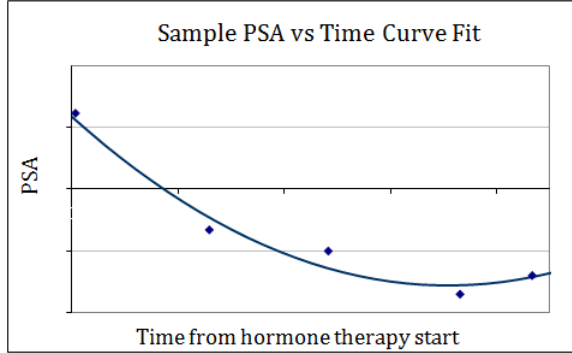
The PSA progression of patient  $i$  as a function of the time  $t$  elapsed since the patient started his hormone treatment (in days) can be modeled by a log quadratic curve:

$$Y_{i,t} = \ln(\text{PSA}_{i,t}) = \alpha_i + \beta_i t + \gamma_i t^2 + \varepsilon_i \quad \varepsilon_i \sim N(0, V_i) \quad (1)$$

Equation (1) can be justified in two ways: theoretically (see Appendix B) and empirically (Lavieri et al. 2009). Assume that the size of the tumor, represented by the number of tumor cells, is proportional to the levels of PSA in the blood. We model two kinds of cells: cells that can divide in the absence of androgens and cells that require androgens to reproduce. When hormone therapy is given, this will decrease the levels of androgens and/or their effectiveness: the cell death rate is higher than the cell division rate, and a net decrease in the size of the tumor is observed. Over time, cells become resistant to the hormone therapy, and as cells increase their division rate, the tumor might start to grow again.

A typical curve appears in Figure 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 (Lavieri et al. 2009).

**Figure 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.

## 2.2 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 (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:

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

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:

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

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).

### 2.3 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:

$$tnadir_i = -\beta_i / (2\gamma_i) \quad (4)$$

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:


$$tnadir_i = \begin{cases} \min(-\beta_i / (2\gamma_i), 240) & \gamma_i * \beta_i < 0 \\ 240 & \text{otherwise} \end{cases} \quad (5)$$

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

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 (see Appendix B). Therefore,  $\beta_i$  and  $\gamma_i$  are negatively correlated. As can be seen in Figure 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. 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  $\begin{pmatrix} r_{22_i} & 0 \\ 0 & r_{33_i} \end{pmatrix}$  and  $\begin{pmatrix} r_{22_i} & r_{23_i} \\ r_{32_i} & r_{33_i} \end{pmatrix}$  respectively.

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

$$F(t_{nadir}) = \frac{1}{\sqrt{2\pi}} \left( \int_{-\infty}^{\frac{(\hat{\beta}_i + 2\hat{\gamma}_i t)}{\left(2\sqrt{r_{33_i}}\sqrt{r_{22_i}}\sqrt{\frac{t^2}{r_{22_i}} + \frac{r_{32_i}t}{r_{22_i}r_{33_i}} + \frac{1}{4r_{33_i}}}\right)}} \exp\{-y^2/2\} dy \right) \quad (6)$$

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 (6):

$$|F(t_{nadir}) - F^*(t_{nadir})| \leq 1 - \Phi\left(r_{32} / \sqrt{r_{33} * r_{22}}\right) \quad (7)$$

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  $\begin{pmatrix} r_{22} & r_{23} \\ r_{32} & r_{33} \end{pmatrix}$ . Using equation (5), we estimated the time of nadir for each observation. By discretizing time, we estimated the proportion of times (out of 10000) that the simulated nadir fell into each time category. As can be seen in Appendix D, the two methods of estimating the period of maximum probability gave similar results, suggesting that this approximation is appropriate.

#### **2.4 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, Sjöström and Nygaard 2005), monitoring glucose levels in the blood (Parker, Doyle and Peppas 1999, Knobbe and Buckingham 2005) and tracking the position of the tumour in patients receiving radiation therapy (Rehbinder, Forsgren

and Löf 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 ( $\ln(\text{PSA})$ ), 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:

$$Y_{i,t} = F_{i,t}^T \theta_{i,t} + \varepsilon_{i,t} \quad \varepsilon_{i,t} \sim N(0, V_{i,t}) \quad (8)$$

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:

$$\theta_{i,t} = \theta_{i,t-1} + \omega_{i,t} \quad \omega_{i,t} \sim N(0, W_{i,t}) \quad (9)$$

where  $\omega_{i,t}$  is the  $3 \times 1$  vector of the parameters' errors with mean zero and covariance matrix  $W_{i,t}$ . Given our assumption that the 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 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:

$$\hat{\theta}_{i,t|t} = \hat{\theta}_{i,t|t-1} + R_{i,t|t-1} F_{i,t} Q_{i,t}^{-1} [Y_{i,t} - F_{i,t}^T \hat{\theta}_{i,t|t-1}] \quad (10)$$

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

$$R_{i,t|t} = R_{i,t|t-1} - R_{i,t|t-1} F_{i,t} Q_{i,t}^{-1} F_{i,t}^T R_{i,t|t-1} \quad (11)$$

$$Q_{i,t} = F_{i,t}^T R_{i,t|t-1} F_{i,t} + V_{i,t} \quad (12)$$

From equation (9),

$$\hat{\theta}_{i,t|t-1} = \hat{\theta}_{i,t-1|t-1} \quad (13)$$

and

$$R_{i,t|t-1} = R_{i,t-1|t-1} + W_{i,t} \quad (14)$$

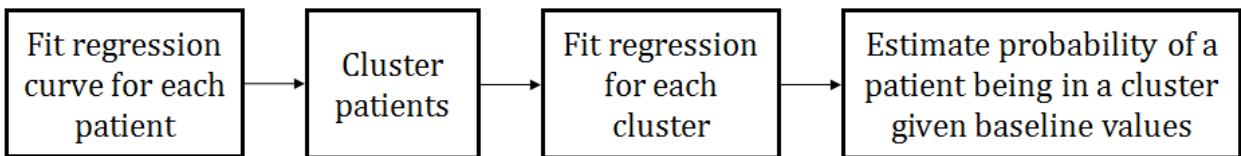
Furthermore, observe that  $Y_{i,t}$  enters only in equation (10). Since the covariance of the parameter estimates doesn't depend on the data obtained each period, it can be computed off-line based on the prior information.

## 2.5 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 4 shows how clustering can be used to improve our model estimates.


As in Section 2.2, we first fit regression curve (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  $minPSA_i$  give results that are easiest to interpret.

**Figure 4: Main Steps in Clustering the Prior Distribution.**



As can be observed in Figure 5, the population is not homogeneous: there appears to be subgroups within the population.

**Figure 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, 2002, 2006). This method combines agglomerate hierarchical clustering with the EM algorithm to maximize the likelihood that observation  $x_i$  belongs to group  $j$ :

$$L(\rho_1, \dots, \rho_M; P(1), \dots, P(M)) = \prod_{i=1}^n \sum_{j=1}^M P(j) f_j(x_i | \rho_j) \quad (15)$$

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 (15) 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 clustering partition is 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 (2) and (3) can be rewritten as:

$$\bar{\theta}_{k,j} = \sum_{i=1}^n c_{i,j} \hat{\theta}_{k,i,j} / \sum_{i=1}^n c_{i,j} \quad \text{for } k = 1, 2, 3 \quad (16)$$

and

$$\bar{\theta}_{k,j}^w \quad (17)$$

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

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. 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.

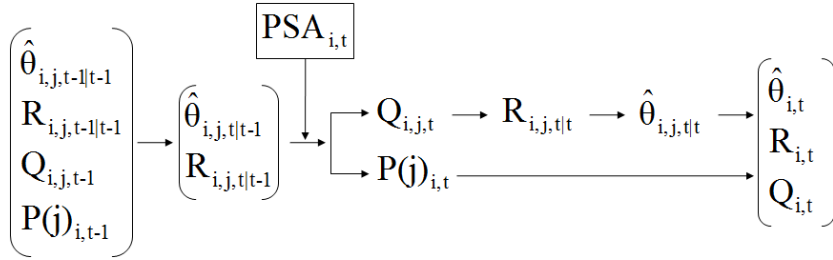
We can proceed in two ways. One way is to assign patients initially to a cluster and only update the parameters for that cluster as new information becomes available. Another possibility would be to update the probability that patient  $i$  belongs to cluster  $j$  at time  $t$  ( $P(j)_{i,t}$ ) as new 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:

$$P(j)_{i,t} = \frac{P(j)_{i,t-1} * f\left(PSA_{i,t} | (F_{i,t}^T \hat{\theta}_{i,j,t|t-1}; Q_{i,j,t-1})\right)}{\sum_{j=1}^3 \left( P(j)_{i,t-1} * f\left(PSA_{i,t} | (F_{i,t}^T \hat{\theta}_{i,j,t|t-1}; Q_{i,j,t-1})\right) \right)} \quad (18)$$

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

Equation (18) 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 6 summarizes the parameter updating process.

**Figure 6. Steps involved in the curve parameter updating.**



### 3. 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.

### 3.1 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 7 shows a scatterplot of the regression parameters obtained for all patients. As can be seen in Figure 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.

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 regressions parameters  $\hat{\alpha}_i, \hat{\beta}_i$  and  $\hat{\gamma}_i$ , the estimated

minimum PSA value (estimated in a similar way as *tnadir*) 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 quite 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 (Lavieri et al. 2009).

Figure 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.

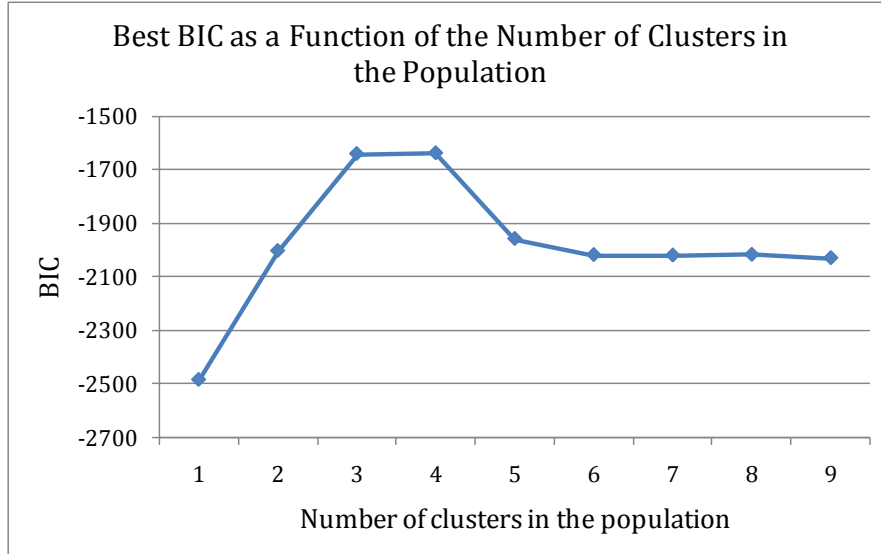


*Patients have been clustered based on  $tnadir_i$  and  $minPSA_i$ .*

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 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 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 1 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.

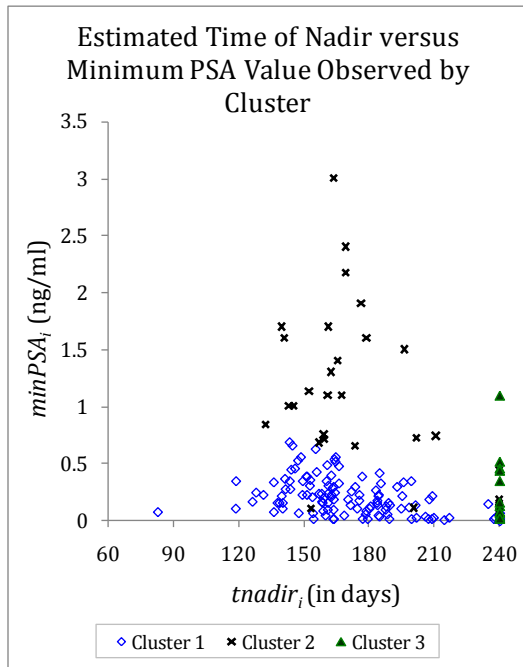
**Table 1. Average cluster probabilities given most likely cluster membership.**

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

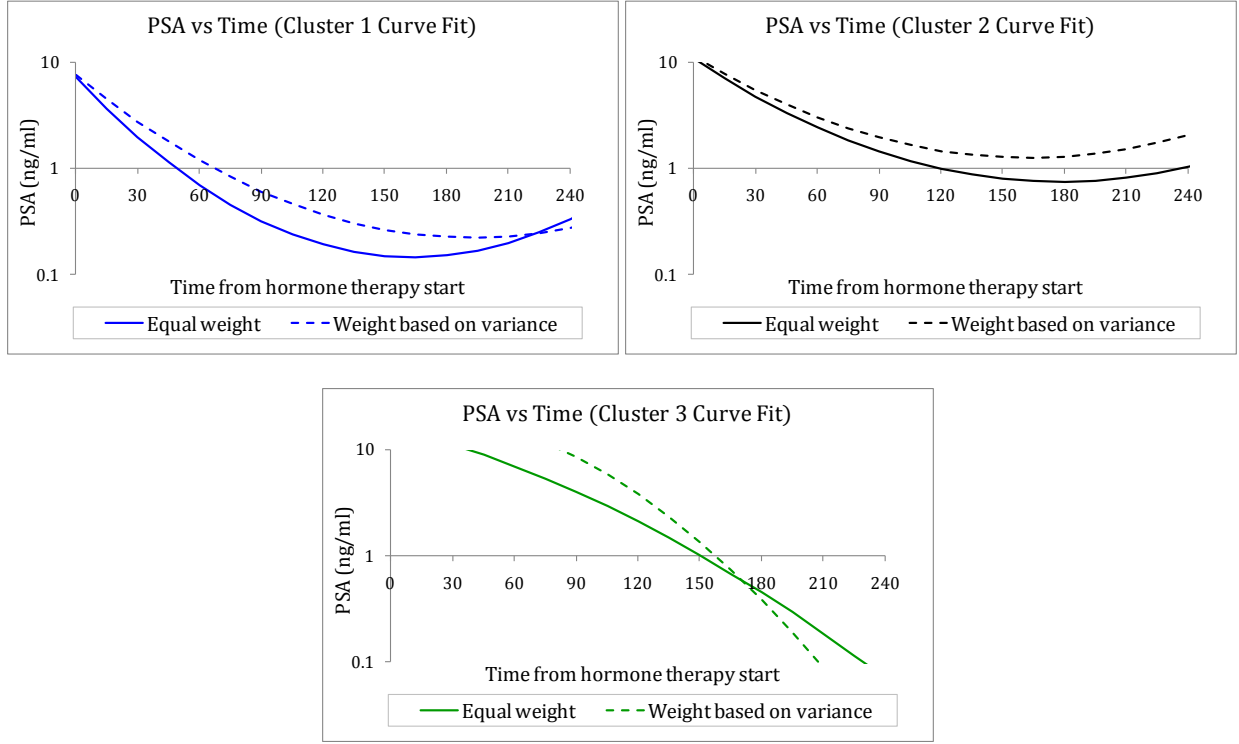


	Cluster 1	Cluster 2	Cluster 3
<i>tnadir</i> (mean, in months)	5.5	8	>> 8
<i>minPSA</i> (mean [variance])	.21[.03]	1.03[.56]	.26[.09]
Proportion of patients in cluster	72%	20%	8%

From Figure 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 the 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 the clinicians.

Using equations (16) and (17), we estimated the parameters within each cluster. The results obtained are presented in Figure 10. The impact on policy decisions of using either equation to estimate the initial parameters within each cluster is shown in Appendix D.

Figure 10. Regression curves based on prior means obtained by weighting parameters within each cluster.



Parameters are weighted based on Equation (16) (equal weight) or on Equation (17) (weight based on variance).

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

$$P(1)_{i,0} = \frac{\exp(2.3 - 0.16 * PSA_{i,0})}{1 + \exp(2.3 - 0.16 * PSA_{i,0}) + \exp(1.16 - 0.06 * PSA_{i,0})} \quad (19)$$

$$P(2)_{i,0} = \frac{\exp(1.16 - 0.06 * PSA_{i,0})}{1 + \exp(2.3 - 0.16 * PSA_{i,0}) + \exp(1.16 - 0.06 * PSA_{i,0})} \quad (20)$$

$$P(3)_{i,0} = 1 - P(1)_{i,0} - P(2)_{i,0} \quad (21)$$

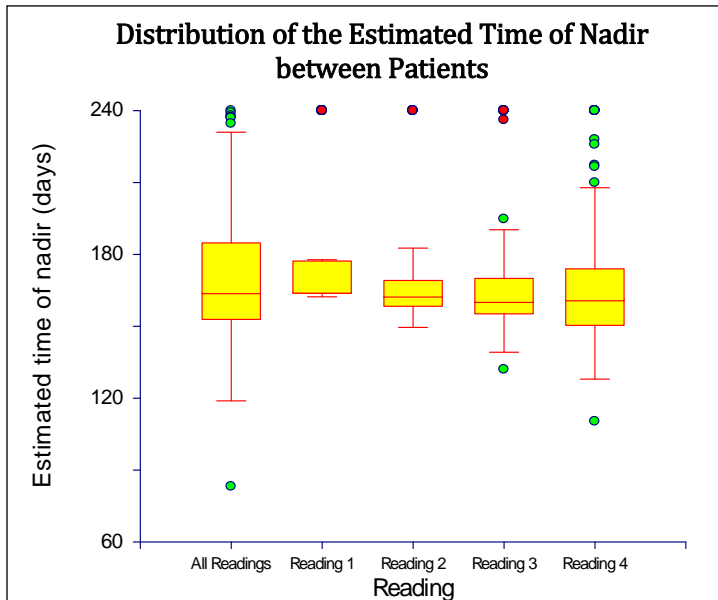
$P(j)_{i,0}$  represents the probability that a patient  $i$  belongs to cluster  $j$  at the beginning of the neoadjuvant hormonal treatment (e.g. at time 0).  $PSA_{i,0}$  is the PSA value for patient  $i$  at time 0.

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 become available for a given patient, we used the Kalman filter

to update the regression parameter estimates. Based on Equation (5), we then estimated the time of nadir.

Figure 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).

**Figure 11. Distribution of the estimated time of nadir between 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 (5) 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 (1) retrospectively for each patient given all PSA readings from the start of the hormone treatment to the start of the RT.*

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  $tnadir_i$ .

We continue to update the probability a patient is in each cluster using Equation (18) 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 12 summarizes how clustering classification changes over time.

As the patient’s curve parameters are updated, Clusters 1 and 2 become quite 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 both costly and has associated side effects.

**Figure 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 (18). PSA readings are taken every 2 months on average.*

### 3.2 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. Input from our clinical collaborators was crucial in devising decision rules that would be acceptable in practice. Each month is assumed to consist of 30 days. Since we can't 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 (5), based on the regression parameters calculated retrospectively given all PSA readings for that patient. With the exception of the current protocol, all policies used 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 are receiving 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
- PSA < 0.05 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 eight month, it is assumed that all patients receive RT at the eight 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.

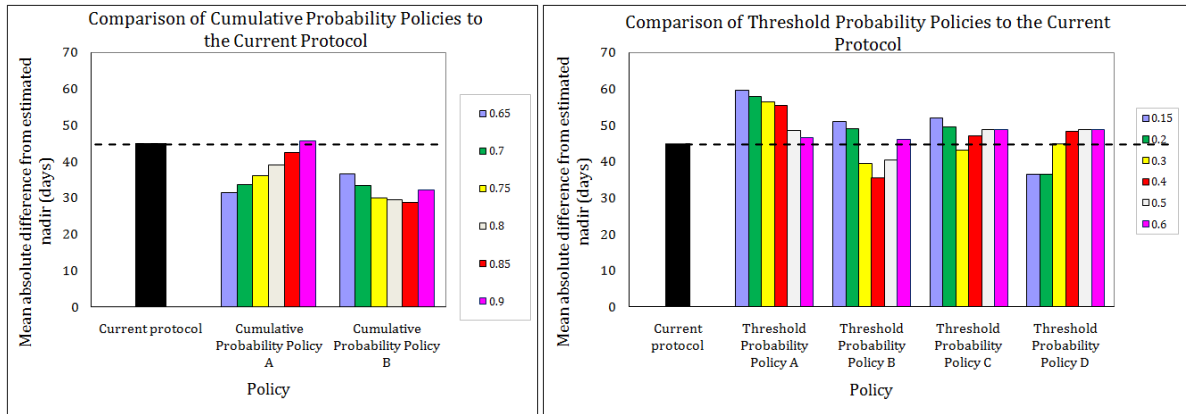
- Threshold probability policy D: This policy starts the radiotherapy treatment of the patient if the patient's probability of reaching the nadir from two weeks before the PSA reading until two weeks after that reading is higher than a given threshold.

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

The policies are compared in Figure 13 based on the mean absolute difference between the policy's treatment time and the patient's estimated time of nadir (see Table D.1 in Appendix D for further detail).

From Figure 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.

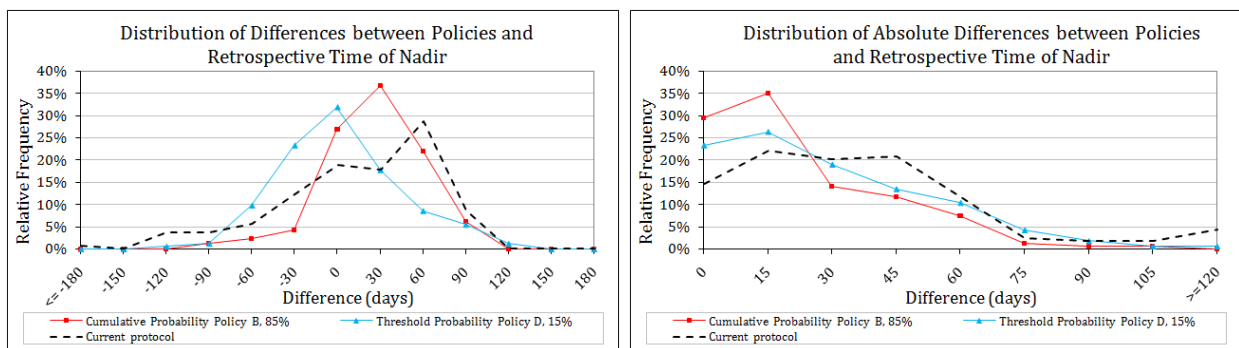
**Figure 13. Comparison of the mean absolute difference between the RT start time based on each policy and the estimated time of nadir for each patient in the population.**



The estimated time of nadir is based on Equations (1) and (5) using all PSA readings for that patient (not achievable). 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 14.

**Figure 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 (1) and (5) 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 a wider 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 to the 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 increased disease progression, increased risk of cells becoming resistant to the 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 (Lavieri et. al 2009). 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 in the time of treatment to the time of treatment under their current protocol is 74 and 103 days respectively.

Various modifications to the model are discussed in Appendix D. Modifications considered include:

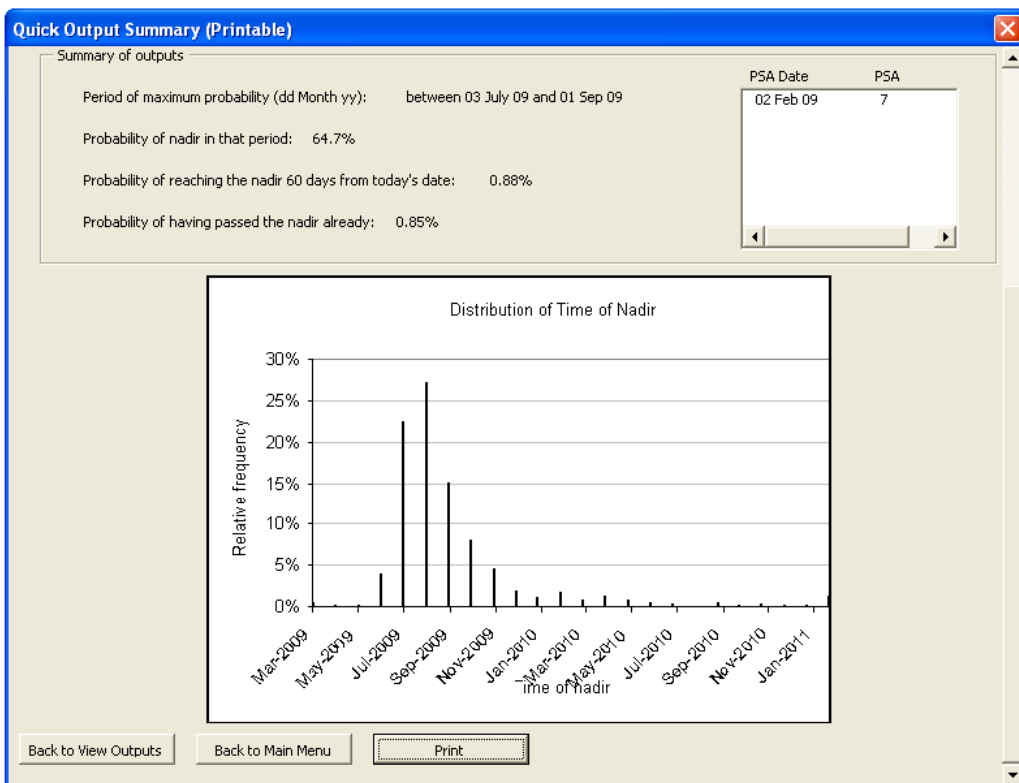
- Estimating the nadir probability using simulation instead of the integral approximation described in equation (6)
- Using a weighted probability of the nadir within all clusters in each period rather than the probability of nadir of the cluster for which  $P(j)_{i,t}$  is maximized.
- Using Equation (16) rather than equation (17) to weight the prior estimates of the parameters within each cluster.
- Incorporating measurement error.

As can be seen in Appendix D, none of these changes improves our solutions significantly, suggesting that our initial assumptions are appropriate.

In addition, we compare our results to the results obtained if clustering was to be omitted from our analysis. That is, if we assume that all patients in the population belong to a single cluster (see Appendix D). 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 15.

**Figure 15. Sample output obtained from the tool.**



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.

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.

#### **4. Conclusion 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 radiotherapy. By using a threshold to decide whether to start RT, we are able to identify earlier when the nadir is 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 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 are 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 covariance 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 has 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.

## Acknowledgements

This work was partially supported by 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. The authors thank William J. Morris for provision of data based and manuscript revision. Any errors in the data analysis or interpretations are the first three authors'. The authors thank Dr. Steven Shechter and members of the CIHR Team in Operations Research for Improved Cancer Care for their helpful comments.

## Appendix A. Model Notation

We summarize the notation used in our model in Table A.1.

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
$\bar{\theta}_j = (\bar{\alpha}_j, \bar{\beta}_j, \bar{\gamma}_j)^T$	
$\bar{\theta}_j^w = (\bar{\alpha}_j^w, \bar{\beta}_j^w, \bar{\gamma}_j^w)^T$	
$\hat{\theta}_i = (\hat{\alpha}_i, \hat{\beta}_i, \hat{\gamma}_i)^T$	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 $j$

Table A.1. Model notation.

## Appendix B. PSA Dynamics

Here we derive expression (1). Assume that the size of the tumor  $N_i(t)$ , as 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 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:

$$dN_i(t)/dt = (g_i - a_i)N_i(t) \quad (\text{B.1})$$

By solving (B.1),  $N_i(t)$  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 and/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

$$dZ_i(t)/dt = g_i tZ_i(t) - a_iZ_i(t) \quad (\text{B.2})$$

and

$$dX_i(t)/dt = -a_iX_i(t) \quad (\text{B.3})$$

If we let  $m_i$  be 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  $Z_i(t) = m_iN_i(t)$ . From equations (B.2) and (B.3), the net growth rate can be rewritten as:

$$dN_i(t)/dt = g_i m_i t N_i(t) - a_i N_i(t) \quad (\text{B.4})$$

By solving (A.4), we obtain  $N_i(t) = C_i \exp(1/2g_i m_i t^2 - a_i t)$ . Under the assumption of normality of the errors of the PSA readings, the PSA dynamics can therefore be modeled as:

$$Y_i = \ln(\text{PSA}_i) = \alpha_i + \beta_i t + \gamma_i t^2 + \varepsilon_i \quad \varepsilon_i \sim N(0, V_i) \quad (\text{B.5})$$

### Appendix C. Derivation of the Approximation of the Distribution of the Time of Nadir

Following results by (Hinkley, 1969), we proceed to show that the distribution of the time of nadir converges to (6) as  $P(\gamma_i > 0) \rightarrow 1$ .

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

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

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

and

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

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

$$\frac{1}{2\pi\sqrt{1-g(t)^2}} \left( \int_{h(t)}^{\infty} \int_{\frac{-\hat{\gamma}_i}{\sqrt{r_{33_i}}}}^{\infty} \exp\left\{\frac{-x^2 + 2g(t)xy - y^2}{2(1-g(t)^2)}\right\} dx dy + \int_{-h(t)}^{\infty} \int_{\frac{\hat{\gamma}_i}{\sqrt{r_{33_i}}}}^{\infty} \exp\left\{\frac{-x^2 + 2g(t)xy - y^2}{2(1-g(t)^2)}\right\} dx dy \right) \quad (\text{C.3})$$

Letting  $P(\gamma_i > 0) \rightarrow 1$  (or  $\hat{\gamma}_i/\sqrt{r_{33_i}} \rightarrow \infty$ ), the second term in equation (C.3) approaches 0. Given

that  $\int_{-\infty}^{\infty} \exp\{-ax^2 - 2bx\} dx = \sqrt{\pi/a} \exp\{b^2/a\}$ , the first term in equation (C.3) 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:

$$\frac{1}{\sqrt{2\pi}} \left( \int_{h(t)}^{\infty} \exp\{-y^2/2\} dy \right) \quad (\text{C.4})$$

Which is equivalent to:

$$\frac{1}{\sqrt{2\pi}} \left( \int_{-\infty}^{-h(t)} \exp\{-y^2/2\} dy \right) \quad (\text{C.5})$$

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:

$$\frac{1}{\sqrt{2\pi}} \left( \int_{-\infty}^{h(t)} \exp\{-y^2/2\} dy \right) \quad (\text{C.6})$$

## Appendix D. Model Comparisons

Here we compare the following modeling approaches based on the mean and maximum absolute difference between the policy treatment time and the patient's estimated time of nadir:

**Baseline:** This is the modeling approach that is followed in the data analysis and implementation section. 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 (6) to estimate the nadir probabilities given that the patient belongs to the cluster with highest probability.
- Uses equation (16) to weight the prior estimates of the parameters within each cluster.
- Does not incorporate measurement error.

**Simulation:** This modeling approach calculates the probability of the nadir by using simulation instead of the integral approximation described in equation (6). 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 (17) is used instead of equation (16).

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

**No clusters:** This approach assumes that no clusters are present in the population by letting all patients be part of a single cluster.

Policies are compared in Table D.1

Policy	Threshold Used	Baseline		Simulation		Weight cluster probabilities		Weight priors		Include measurement 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

Table D.1 Model comparisons based on the mean absolute difference and maximum absolute difference between each policy's RT start time and the estimated  $tnadir_i$  for all patients. The minimum mean and maximum difference within each policy and modeling approach is highlighted.

## References

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

- American Cancer Society. 2008. Detailed guide: prostate cancer. Retrieved June 30, 2009, [www.cancer.org/docroot/CRI/CRI\\_2\\_3x.asp?dt=36](http://www.cancer.org/docroot/CRI/CRI_2_3x.asp?dt=36).
- Agresti, A. 1996. *An introduction to categorical data analysis*. John Wiley & Sons, New York.
- BC Cancer Agency. 2009. Prostate cancer: BC Cancer Agency. Retrieved June 30, 2009, [www.bccancer.bc.ca](http://www.bccancer.bc.ca).
- Celeux, G., G. Soromenho. 1996. An entropy criterion for assessing the number of clusters in a mixture model. *Journal of Classification* **13** 195-212.
- 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.
- Fieller, E. C. 1932. The distribution of the index in a normal bivariate population. *Biometrika* **24** 428-440.
- Fraley, C., 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., A. E. Raftery. 2002. Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association* **97** 611-631.
- Fraley, C., A. E. Raftery. 2006. MCLUST version 3 for R: normal mixture modeling and model-based clustering. Technical Report No. 504, Department of Statistics, University of Washington.
- Gleave, M. E., S. La Bianca, 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., 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.
- Harvey, A. 1991. Forecasting structural time series models and the Kalman filter. Cambridge University Press, Cambridge 100-167.
- Hinkley, D. V. 1969. On the ratio of two correlated normal random variables. *Biometrika* 56(3) 635-639.
- Heymann, J. J., M. C. Benson, K. M. O'Toole, B. Malyszko, R. Brody, D. Vecchio, P. B. Schiff, M. M. Mansukhani, 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., G. E. Hanks. 2000. Hormonal therapy combined with radiation or surgery in treatment of locally advanced non-metastatic prostate cancer. *Comprehensive Textbook of Genitourinary Oncology*, eds. N. J. Vogelzang, P. T. Scardino, W. U. Shipley, D. S. Coffey, B. J. Miles. Lippincott Williams & Wilkins, Philadelphia 782-789.
- Knobbe, E. J., B. Buckingham. 2005. The extended Kalman filter for continuous glucose monitoring. *Diabetes Technology & Therapeutics*. 7(1) 15-27.
- Kristiansen, N. K., S. -O. Sjöström, H. Nygaard. 2005. Urinary bladder volume tracking using a Kalman filter. *Medical and Biological Engineering and Computing*. 43(3) 331-334.
- Lavieri, M. S., M. L. Puterman, S. Tyldesley, et al. 2009. A new approach for deciding when to start radiation therapy in prostate cancer patients receiving neoadjuvant hormone therapy. manuscript in preparation.

- Makarov, D. V., B. J. Trock, E. B. Humphreys, L. A. Mangold, P. C. Walsh, J. I. Epstein, 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.
- 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. [www.cancer.gov/search/ViewClinicalTrials.aspx?cdrid=453036&version=HealthProfessional&protocolsearchid=6093623](http://www.cancer.gov/search/ViewClinicalTrials.aspx?cdrid=453036&version=HealthProfessional&protocolsearchid=6093623).
- National Cancer Institute. 2009. Cancer Statistics Home Page – National Cancer Institute. Retrieved November 12, 2009. [www.cancer.gov/statistics/](http://www.cancer.gov/statistics/).
- 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.
- 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. , B. Carter. 2000. Use of prostate-specific antigen density and velocity. *Comprehensive Textbook of Genitourinary Oncology*, eds. N. J. Vogelzang, P. T. Scardino, W. U. Shipley, D. S. Coffey, B. J. Miles. Lippincott Williams & Wilkins, Philadelphia 650-654.

- Rehbinder, H., C. Forsgren, J. Löf. 2004. Adaptive radiation therapy for compensation of errors in patient setup and treatment delivery. *Medical Physics*. **31**(12) 3363-3371.
- Sandblom, G., L. Holmberg, J. E. Damber, J. Hugosson, J. E. Johansson, R. Lundgren, E. Mattsson, J. Nilsson, E. Varenhorst. 2002. Prostate-specific antigen as surrogate for characterizing prostate cancer subgroups. *Scandinavian Journal of Urology & Nephrology*. **36**(2) 106-112.
- Shechter, S. M., M. D. Bailey, A. J. Schaefer, M. S. Roberts. 2008. The optimal time to initiate HIV therapy under ordered health states. *Operations Research* **56**(1) 20-33.
- Zhu, H., R. Rohwer. 1996. Bayesian regression filters and the issue of priors. *Neural Computing & Applications*. **4**(3) 130-142.