Published in final edited form as:

Health Care Manag Sci. 2012 December ; 15(4): 293–309. doi:10.1007/s10729-012-9195-x.

Simulation optimization of PSA-threshold based prostate cancer screening policies

Daniel J. Underwood,

Edward P. Fitts Department of Industrial & Systems Engineering, North Carolina State University, Raleigh, NC 27695, USA

Jingyu Zhang,

Philips Research North America, Briarcliff Manor, NY 10510, USA

Brian T. Denton,

Edward P. Fitts Department of Industrial & Systems Engineering, North Carolina State University, Raleigh, NC 27695, USA

Nilay D. Shah, and

Division of Health Care Policy and Research, Mayo Clinic, Rochester, MN 55905, USA

Brant A. Inman

Division of Urology, Duke University Medical Center, Durham, NC 27710, USA

Daniel J. Underwood: daniel.underwood@ncsu.edu; Jingyu Zhang: jingyu.zhang@philips.com; Nilay D. Shah: Shah.Nilay@mayo.edu; Brant A. Inman: brant.inman@duke.edu

Abstract

We describe a simulation optimization method to design PSA screening policies based on expected quality adjusted life years (QALYs). Our method integrates a simulation model in a genetic algorithm which uses a probabilistic method for selection of the best policy. We present computational results about the efficiency of our algorithm. The best policy generated by our algorithm is compared to previously recommended screening policies. Using the policies determined by our model, we present evidence that patients should be screened more aggressively but for a shorter length of time than previously published guidelines recommend.

Keywords

Prostate cancer screening; Simulation optimization; Genetic algorithm; Ranking and selection

1 Introduction

Prostate cancer is one of the most common cancers affecting American men. Statistics show for males in the U.S. during 2009 that prostate cancer represented 25% of new cancers and caused 9% of deaths from cancer (excluding skin cancers) [33]. The evidence for benefits of prostate-specific antigen (PSA)-based prostate cancer screening is unclear, and the issue of whether and how to screen is heavily debated among healthcare professionals and policymakers. Randomized controlled trials (RCTs) have recently been undertaken in the U.S. and in Europe to investigate the effects of different screening policies [3, 52]. However, RCTs are time-consuming, costly, and subject to problems such as limited sample size,

selection biases, and patient adherence. Adding to the confusion, results from the above referenced RCTs have produced conflicting results. In the European study, PSA-based screening reduced the rate of death from prostate cancer by 20%; whereas, in the U.S. study, the rate of death from prostate cancer for the screen group was not significantly different from the rate for the control group.

Clinicians screen for prostate cancer by monitoring the serum prostate-specific antigen (PSA) level. PSA, a protein produced by prostate cells, is a continuous biomarker that is associated with prostate cancer. High PSA levels often point to the existence of cancerous cells in the prostate but may also point to benign enlargement of the prostate gland, a physiologic process that occurs with age. Due to the presence of other potential causes for elevated PSA, PSA screening is an imperfect test and can produce false positive results. Furthermore, because the likelihood of dying from causes other than prostate cancer increases dramatically with age, and because prostate cancer is a slow growing cancer, PSA screening loses its benefits in older patients. The major decisions related to the PSA-screening debate include how often to screen an individual, when to discontinue screening, and at what PSA threshold is prostate biopsy necessary.

Prostate cancer screening decisions are difficult. Early detection and treatment of prostate cancer can add decades to a patient's life. However, because of the high risk of false positives and the slow-growing nature of prostate cancer, physicians may be hesitant to recommend prostate biopsy for fear of overdiagnosis and overtreatment. Additionally, patients themselves can be reluctant to undergo biopsy due to the invasive nature of the procedure.

In this article we present a simulation-optimization model that combines a simulation model with a genetic algorithm (GA) and a probabilistic ranking and selection method to design a PSA-threshold based screening policy. Our PSA sampling method is based on a large longitudinal data set of patients at Mayo Clinic in Rochester, MN. We discuss the performance of the GA and present findings derived from comparing previously published PSA screening policies to a policy designed by our simulation-optimization model.

The remainder of this article is organized as follows: we provide background on prostate cancer in Section 2; we review the relevant literature in Section 3; we present the simulation-optimization model in Section 4; we report on results from the simulation-optimization model in Section 5; and, finally, we make concluding remarks in Section 6.

2 Background on prostate cancer

Physicians commonly obtain a serum PSA sample to screen for prostate cancer [45]. PSA testing consists of drawing a sample of blood and analysing it for its total PSA content. The result of this analysis is a quantity measured in ng/mL of PSA in the blood. If the amount exceeds some threshold, typically 4 ng/mL, then the PSA test is interpreted by the physician as positive. The goal of PSA screening is to detect the prostate cancer while it is localized to the prostate gland so that treatments may be applied that prevent cancer from spreading to other organs.

There are different screening policies in practice across the world [3, 52, 57]. Each policy suggests a frequency for screening and the threshold level of PSA that, when reached, constitutes a positive PSA test result and signals a need for prostate biopsy. PSA is measured on a continuous scale and must be interpreted in light of specified threshold values. Different thresholds result in different sensitivities and specificities, which can be shown using Receiver Operating Characteristic (ROC) curves. The ROC curve for PSA testing, from Zhang et al. [64], is shown in Fig. 1. Studies on the value of PSA screening for

prostate-cancer detection have produced mixed results. The first report from the Prostate, Lung, Colorectal, and Ovarian (PLCO) Cancer Screening Trial on prostate-cancer mortality indicates that PSA screening does *not* effectively reduce the risk of death from prostate cancer [3]. Results from the European Randomized Study of Screening for Prostate Cancer (ERSPC), on the other hand, indicate that PSA screening *does* reduce the risk of death from prostate cancer [52]. It should be noted that 52% of the non-screening group actually underwent screening in the PLCO study [53], which has been cited as a significant source of bias toward finding no benefit from screening [18]. Also, in the ERSPC study the results were highly biased by a single testing center in Sweden that demonstrated a more favorable effect on reduction in prostate cancer mortality than other centers, and when the results from the center in Sweden were removed, significantly less overall benefit from PSA screening was found by the ERSPC study [58].

When a physician suspects a patient has prostate cancer due to high serum PSA levels, the patient is usually referred for biopsy. During prostate biopsy, 16–18 gauge needles are inserted through the rectum into the prostate gland and cores of tissue are removed for further inspection under a microscope. The ideal number of cores to remove is not known but most physicians remove between 12 and 20 cores on a standard biopsy [43]. The biopsy process is guided by transrectal ultrasound, which requires inserting a probe into the rectum so that the physician can image the prostate area and accurately place the needles into specific anatomic locations within the prostate gland. Patients generally view this process as invasive and uncomfortable, and a definite health utility decrement is associated with prostate biopsy. Furthermore, prostate biopsy is an imperfect test for determining the presence of prostate cancer; Haas et al. [31] estimate that the prostate cancer detection rate for prostate biopsy is about 0.8. For a detailed discussion of prostate biopsy standards, see Djavan and Margreiter [17].

Because prostate biopsy involves traversing the rectal wall with needles, there is risk of bleeding and infection [61]. The treatment options following a positive biopsy include, most commonly, radical prostatectomy (surgical removal of the prostate gland) [60] and, less commonly, brachytherapy, external beam radiotherapy, or active surveillance. All prostate cancer treatment options can have serious negative side-effects, such as urinary incontinence and sexual dysfunction [48]. Moreover, because the all-cause mortality rate is high for the age at which many patients undergo prostate biopsy, patients often die of causes other than prostate cancer, thus not benefiting from the prostate biopsy or treatment. For these reasons, researchers, patients, and practitioners are strongly motivated to study and improve prostate cancer screening policies.

3 Literature review

In this section, we review the relevant literature on modeling prostate cancer screening; we review published guidelines for prostate cancer screening; we review the relevant literature on simulation-optimization and ranking-and-selection; and, finally, we explain this article's contribution to the literature.

3.1 Modeling prostate cancer screening

Researchers from the Cancer Intervention and Surveillance Modeling Network of the National Cancer Institute, the *Prostate Working Group*, have studied, developed, and compared models of the natural progression of prostate cancer for over 10 years [41]. Draisma et al. [19] modeled the development of prostate cancer up to the point of diagnosis as a Markov model with states defined by cancer stage and cancer differentiation grade, and produced evidence that PSA screening at intervals larger than 1 year is effective. Etzioni et al. [23] examined the frequency at which healthy males should undergo PSA testing by

using a cancer-stage based Markov natural history model combined with a PSA screening simulation model, and found that biannual screening is a cost-effective alternative to annual screening. Building on [23], Etzioni et al. [24] examined multiple birth cohorts between 1980 and 2000 and found that PSA screening accounts for 80% of the decline in distant stage incidence since the pre-PSA screening era.

Ross et al. [47] compared various PSA screening policies using a Monte-Carlo simulation model based on a Markov model of the natural progression of prostate cancer; the authors concluded that the standard U.S. policy does not prevent more prostate cancer deaths than a policy in which screening begins earlier and occurs less frequently. This study also evaluated different screening strategies by varying the PSA threshold for biopsy, the frequency with which to screen, and the start age for screening; these strategies were evaluated in terms of the number of prostate cancer deaths prevented and the number of PSA tests administered.

Zhang et al. [63] studied optimal screening policies by solving a partially-observable Markov decision process (POMDP) with the goal of maximizing expected quality-adjusted life years. Zhang et al. [63] computed the optimal policy on the belief (probability) of having prostate cancer, and PSA results were treated as information used to update the belief; whereas we generate policies based on PSA thresholds using simulation-optimization, and our results are easier to understand and to use by physicians.

Researchers have also used quantitative models to study other forms of cancer. For example, Maillart et al. [37] assessed the value and frequency of mammography screening by formulating a partially observed Markov chain to determine an efficient set of screening policies, and found that, regardless of frequency, screening should start relatively early in life and continue until relatively late in life. Using a differential equation model of the natural history of cervical cancer, Gustafsson and Adami [30] implemented a variant of the Nelder–Mead Simplex algorithm to find the set of ages when screening for cervical cancer is most effective, and computed optimal schedules for screening dependent upon the number of times to screen.

3.2 Prostate cancer screening guidelines

Authorities in the U.S. are not unanimous regarding PSA screening policy recommendations. The American Urological Association and the American Cancer Society recommend annual PSA tests for men above age 50, and starting even earlier for men in high risk categories [1, 2]. The National Comprehensive Cancer Network recommended an algorithm to determine the best screening policy based on various risk factors, such as a family history of prostate cancer [42]. The American College of Physicians and the American College of Preventative Medicine recommended against routine PSA screening [25]. The U.S. Preventative Services Task Force (USPSTF) previously recommended that males should only be screened under the age of 75 [57]. These recommendations have resulted in the *de facto* standard U.S. PSA screening policy of annual PSA tests from age 50 to 75 with threshold 4.0 ng/mL [47]. More recently, however, the USPSTF issued a draft recommendation statement recommending against routine PSA screening for prostate cancer [58].

European countries also differ in their recommended PSA screening policies. Based on the European Randomized Study of Screening for Prostate Cancer, most European countries (that participated in the study) use a threshold of 3.0 ng/mL [52]. However, Finland, Italy, Holland, and Belgium use a threshold of 4.0 ng/mL, like the U.S. [52]. Most European countries use 4-year screening intervals, although Sweden uses 2-year intervals [52].

3.3 Simulation optimization

Simulation optimization is any instance of optimization where the objective function is evaluated by simulating a stochastic system; the simulation is effectively a black box. Fu [26] reviews a host of optimization techniques, including tabu search, simulated annealing, and GAs, which are suited to optimize problems in which the objective function must be approximated by simulation, and discusses the challenges presented by and the growing interest in simulation optimization. Simulation optimization problems are often solved by integrating a simulation model in a GA, as described in Eskandari et al. [22] and Deb et al. [16], especially when the problems can be easily represented using coded strings (e.g. string of PSA thresholds); a coded string is any finite string or vector that fully describes a solution when decoded and that can be mutated and then decoded to fully describe a different solution.

Buchholz and Thümmel [11] present a strategy for solving simulation optimization problems that utilizes a ranking and selection procedure as part of an evolutionary algorithm. In the context of choosing the best among several alternatives, ranking and selection procedures control the probability that the best alternative simulated system is the alternative which is actually selected; in the literature, this probability is referred to as the *probability of correct selection*.

Rinott [46] presented a two-stage procedure to guarantee a specified lower bound on the probability of correct selection. This early procedure was an improvement on the work by Dudewicz and Dalal [20] wherein weighted averages of the observations in each random sample are compared; in the Rinott procedure, ordinary sample means are compared. Both the Rinott method and the Dudewicz and Dalal method are concerned with selecting the single best alternative with some prescribed probability of correct selection. When the objective is to select a set containing the best alternative—rather than selecting the best alternative itself—the ranking and selection procedure is commonly referred to as a subset selection procedure. Koenig and Law [35] presented a two-stage sampling procedure for selecting a specified number of *best* alternatives from a population. More recent work on subset selection procedures is found in Chen [12]. Ranking and selection procedures have also been extended to Pareto optimization [13]. For excellent surveys on research in the field of ranking and selection, see [29, 32, 54].

3.4 Contribution to the literature

Markov processes are a common way of modeling a patient's natural history or progression through health states [19, 23, 24, 47]. We extend this by integrating a GA in a simulation model that is based on a Markov model similar to that of Ross et al. [47]. To our knowledge, this is the first attempt to optimize prostate cancer screening using a GA.

To our knowledge, this is also the first attempt to use quality adjusted life years (QALYs) to systematically investigate various prostate cancer screening policies. Whereas Etzioni et al. [24] examined prostate cancer screening's impact on distant stage incidence rates, and Ross et al. [47] examined the impact of different policies on prostate cancer mortality and the number of PSA tests and biopsies, we examine the impact of different prostate cancer screening policies on quality of life. According to Schröder and WildHagen [51], the effect of screening policies on mortality rates must be balanced against the effect on cost and quality of life. Our research provides important insight into quality-of-life in the prostate cancer screening debate.

Whereas Zhang et al. [63] computed optimal screening policies that dictate a particular clinical action based on the physician's confidence of the patient being in a particular state (i.e. based on the probability that a patient is in a particular state), the policies we generate

are easier to use clinically because they comprise only screening frequency and thresholds for a positive PSA result, which is precisely the form of screening policy used by clinicians.

4 Simulation-optimization model

In an optimization model, one searches for a solution that minimizes or maximizes some function. The same is true in a simulation-optimization model, except that the function evaluations are carried out by repeated simulations. A typical simulation-optimization model develops candidate solutions, and then simulates these candidate solutions, using the evaluated performance of these candidate solutions to inform the creation of new candidate solutions. In this section we describe a method based on the integration of a simulation model, a GA, and a ranking-and-selection procedure to compute a PSA-based screening policy.

This section is organized as follows: we introduce the simulation model in Section 4.1; we present a method for PSA sampling in Section 4.2; we explain how screening policies are represented in Section 4.3; we discuss the reward vector for our Markov model in Section 4.4; we discuss the transition probabilities for our Markov model in Section 4.5; and, finally, we present the optimization model in Section 4.6.

4.1 Simulation model

Our simulation model is based on a non-stationary finite horizon Markov process. The state definitions are the same as or similar to the state definitions in other studies [47, 63]. The states and possible transitions are illustrated in Fig. 2. The labels NC, C, M, and D represent no cancer, cancer not detected, metastasis, and death, respectively. The aggregated cancer states are organ confined (OC), extraprostatic (EP), and lymph node-positive (LN). States OC, EP, and LN are observable cancer stages in which cancer has been treated but has not metastasized. A patient cannot enter the treatment state (T) until entering state C (i.e. until being diagnosed with prostate cancer). After a patient enters state T, he can transition to state M or to state D. The treatment state in Fig. 2 is an aggregate of multiple cancer states. The simplification of aggregating the cancer states into the single treatment state does not result in loss of accuracy, because the reward function for the treatment state is the expected future rewards associated with the underlying Markov reward chain [64].

The transition probabilities among NC, C, T, M, and D are computed as functions of the annual prostate cancer incidence rates, annual all-cause death rates, annual prostate cancer death rates for patients with undetected cancer and for patients undergoing treatment, and the biopsy detection rate.

Rewards are measured in units of QALYs. The reward function is calculated based on approximations from the literature of the QALYs for a patient in the cancer state, the QALYs cost of a prostate biopsy, and the QALYs for a patient after having developed prostate cancer. For discussion on the merits of using QALYs and similar health-outcome measures, see Gold et al. [28].

4.2 PSA sampling

PSA values are sampled from PSA test results based on a sample of 11,872 men from Olmsted County, MN from 1993 to 2005 (a histogram of the PSA test results is shown in Fig. 3). We used observational data in which patient PSA test results were collected at non-uniform intervals, as is often observed in routine clinical practice. Therefore, we used cubic splines to fit a continuous smooth curve of the PSA values for each individual patient. Cubic spline interpolation is a standard method in numerical analysis that uses a cubic polynomial for each interval between consecutive age points at which PSA tests are done. Based on the

spline fit, the patient's PSA value at every age in 6-month increments within the interval during which they received PSA tests, was estimated. This produced a total of 135,737 total test results (actual and estimated) from 51,294 actual test results. After removing test results that follow cancer diagnosis (we focus on early detection only), we are left with 123,865 "observations" to use in our PSA sampling procedure.

In our sampling procedure, each observation is grouped according to its preceding observation from the same patient. Each observation is assigned to an age interval in the set $\mathbb{T} = \{[0, 49), [50, 59), [60, 69), [70, \infty)\}$ that corresponds to the age in years of the patient when that patient's preceding observation was made. Each observation is further assigned to a PSA interval in the set $\mathbb{P} = \{[0, 1), [1, 2), [2, 3), [3, 4), [4, 7), [7, 10), [10, \infty)\}$ that corresponds to the PSA amplitude in ng/mL of that patient's preceding observation. Thus, the observations are placed into bins formed by the Cartesian product $\mathbb{T} \times \mathbb{P}$. This is done separately for patients who do not develop prostate cancer and for patients who do develop prostate cancer. Thus, there are two bin sets: no cancer bins $\{\mathbb{T} \times \mathbb{P}\}^{NC}$ and cancer bins $\{\mathbb{T} \times \mathbb{P}\}^{Ca}$, where Ca denotes prostate cancer at any stage.

To summarize the sampling procedure, we define the following notation: t refers to a decision epoch, and p_t , a_t , and s_t refer respectively to a patient's PSA value, age, and cancer status during epoch t . In the simulation model, the initial PSA value (p_0) for each patient is randomly sampled from the set of all PSA observations recorded for the simulated patient's initial age. Subsequent PSA values are sampled from the bins described above. In particular, for any decision epoch t , the simulated PSA value, p_t , is a function of the patient's age during the previous epoch, a_{t-1} , the patient's PSA value during the previous epoch, p_{t-1} , and the patient's cancer status during the previous epoch, s_{t-1} . More formally, any sampled PSA value after p_0 is defined by $p_t = g(a_{t-1}, p_{t-1}, s_{t-1})$ where the function $g(a, p, s)$ is defined as locating the bin in $\{\mathbb{T} \times \mathbb{P}\}^s$ that contains a and p , and randomly sampling therefrom.

4.3 Screening policies

For each decision epoch, the screening policy defines whether or not to screen and, if the patient is to be screened, the PSA threshold for biopsy. In our model, if the patient's PSA value is greater than the threshold value corresponding to the patient's age, the patient is referred for biopsy; otherwise, no action is taken and the simulation proceeds to the next epoch.

We define the vector of decision variables that defines the policy, K , as

$$K = (x_1, x_2, \dots, x_N),$$

where $x_i \in \{0, 0.5, \dots, 10, \infty\}$ defines a PSA threshold for biopsy referral at age i over N decision epochs separated by 1-year intervals. A threshold value of ∞ is used to denote no biopsy referral, and therefore no PSA test. Additionally, biopsy history may inform the decision to biopsy. We restrict each patient to a single biopsy because patients having prior negative biopsies are screened differently than those who have not. The probability of having cancer given a prior negative biopsy is lower than the probability of having cancer given no prior biopsy [5, 55]. The biopsy technique itself changes for patients who have had a prior negative biopsy, because the physician will extract more samples to reduce the likelihood of missing cancer that may have been missed on prior biopsy [15, 44]. Moreover, prostate cancer that is not found on first biopsy but is found on subsequent biopsy tends to be clinically less significant than prostate cancer found on first biopsy [8, 21, 38]. For these reasons, following an initial biopsy it is assumed that a patient in our model does not

undergo further PSA testing. In other words, the decision process ends and the patient collects a reward defined by their quality-adjusted survival in the absence of screening.

4.4 Rewards

At the beginning of each decision epoch, the patient accumulates a reward, in QALYs, that represents living in his present health state for one additional epoch. We define rewards conditional on the decision to biopsy (B) or to wait (W) as shown below. Definitions of the parameters used in the reward function are shown in Table 1.

$$\begin{aligned} r(\text{NC}, \text{W}) &= 1 & r(\text{NC}, \text{B}) &= 1 - \mu \\ r(\text{C}, \text{W}) &= 1 & r(\text{C}, \text{B}) &= 1 - \mu - \epsilon \\ r(\text{T}, \text{W}) &= 1 - \epsilon & r(\text{T}, \text{B}) &= 1 - \epsilon \\ r(\text{M}, \text{W}) &= 1 - \gamma & r(\text{M}, \text{B}) &= 1 - \gamma \\ r(\text{D}, \text{W}) &= 0 & r(\text{D}, \text{B}) &= 0 \end{aligned}$$

As described in Table 1, we treat the disutility associated with the prostate biopsy procedure as a one-time decrement in QALYs. This is reasonable because prostate biopsy does not usually cause long-term side effects, such as incontinence or sexual dysfunction, which are common complications associated with prostate cancer treatment. A recent study showed that the rates of men who have symptoms of urinary incontinence or erectile dysfunction return to base-line values by 12 weeks after the biopsy procedure was performed [34].

4.5 Transition probabilities

At the end of each period, the patient makes a health-state transition. This transition occurs according to the one-step transition probability matrix below. The parameters of the transition probability matrices are defined in Table 2. The general one-step transition probability matrix, which is identical to that used Zhang et al. [63], is

$$P = \begin{matrix} & \begin{matrix} \text{NC} & \text{C} & \text{T} & \text{M} & \text{D} \end{matrix} \\ \begin{matrix} \text{NC} \\ \text{C} \\ \text{T} \\ \text{M} \\ \text{D} \end{matrix} & \begin{pmatrix} P_{1,1} & P_{1,2} & 0 & 0 & P_{1,5} \\ 0 & P_{2,2} & P_{2,3} & P_{2,4} & P_{2,5} \\ 0 & 0 & P_{3,3} & P_{3,4} & P_{3,5} \\ 0 & 0 & 0 & P_{4,4} & P_{4,5} \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{matrix}$$

Given the decision to biopsy (B), the one-step transition probability matrix is:

$$\begin{aligned}
P_{1,1}^B &= (1 - d_t)(1 - w_t) \\
P_{1,2}^B &= (1 - d_t)w_t \\
P_{1,5}^B &= d_t \\
P_{2,2}^B &= (1 - f)(1 - d_t)(1 - e_t) \\
P_{2,3}^B &= f(1 - b_t)(1 - d_t) \\
P_{2,4}^B &= fb_t(1 - d_t) + e_t(1 - f)(1 - d_t) \\
P_{2,5}^B &= d_t \\
P_{3,3}^B &= (1 - d_t)(1 - b_t) \\
P_{3,4}^B &= b_t(1 - d_t) \\
P_{3,5}^B &= d_t \\
P_{4,4}^B &= (1 - d_t)(1 - z_t) \\
P_{4,5}^B &= d_t + z_t(1 - d_t)
\end{aligned}$$

In the absence of biopsy (W), the one-step transition probability matrix is:

$$\begin{aligned}
P_{1,1}^W &= (1 - d_t)(1 - w_t) \\
P_{1,2}^W &= (1 - d_t)w_t \\
P_{1,5}^W &= d_t \\
P_{2,2}^W &= (1 - d_t)(1 - e_t) \\
P_{2,3}^W &= 0 \\
P_{2,4}^W &= e_t(1 - d_t) \\
P_{2,5}^W &= d_t \\
P_{3,3}^W &= (1 - d_t)(1 - b_t) \\
P_{3,4}^W &= b_t(1 - d_t) \\
P_{3,5}^W &= d_t \\
P_{4,4}^W &= (1 - d_t)(1 - z_t) \\
P_{4,5}^W &= d_t + z_t(1 - d_t)
\end{aligned}$$

After making a health-state transition, the PSA value for the next epoch of the Markov model is sampled (if the patient survives). The simulation proceeds until the patient transitions to the absorbing state D.

4.6 Simulation optimization

The optimal PSA screening policy, K^* , can be defined as

$$K^* = \operatorname{argmax}_{K \in \mathcal{K}} J(K), \quad (1)$$

where

$$J(K) = E_K \left[\sum_{t=t_0}^T r(s, \alpha) \right]$$

is the value function measuring expected QALYs under policy K ; $\mathcal{K} = \langle k_1 \times k_2 \times \dots \times k_N \rangle$, where $k_i \in \{0, 0.5, \dots, 10, \infty\} \forall i$, defines the feasible set of PSA-threshold screening policies; t_0 is a lower bound on the age at which PSA screening begins; T is an upper bound on the age at which PSA screening is discontinued; and $r(s, \alpha)$ is the annual reward function for state $s \in \{\text{NC, C, T, M, D}\}$ and action $\alpha \in \{\text{wait, biopsy}\}$.

In our simulation-optimization approach, the simulator is embedded in a GA and estimates the expected value of a policy. The major steps of the method are illustrated in Fig. 4. The GA is a single-population globally parallel GA; that is, we repeatedly modify a single population, and the evaluation of each individual (also referred to as “chromosome”) in the population is performed in parallel. The population is modified by applying a crossover operator to pairs of selected individuals, and then by applying a mutation operator. In the remainder of this subsection, we describe in detail the major components of the GA; we describe the ranking-and-selection procedure implemented by the GA; and, finally, we present pseudocode for the GA.

Tournament selection—Pairs of individuals are selected from the population by tournament selection. In tournament selection, two individuals are uniformly randomly drawn from the population, and the individual with greater “fitness” is selected to become a parent for reproduction; fitness is a measure of the quality of an individual (solution), which in our problem is measured in expected QALYs. This process repeats until enough parents are selected to create the next generation of individuals. At each round of tournament selection, all individuals are available; in other words, the individuals are randomly drawn *with replacement* from the population. Our implementation of tournament selection is known as *deterministic* tournament selection, because at each pairwise comparison, the individual with greater fitness is selected with probability 1. For details on fitness proportionate selection, an alternative selection operator we tested, see Underwood [56].

The process of mating pairs of individuals (parents) to produce a new individual (child) is accomplished by crossover. The idea behind crossover is to form a new individual that shares features of two existing individuals. In our problem, the two parents are defined by the following vectors of decision variables (each of which constitutes an “individual”):

$$K_1 = (\phi_1, \phi_2, \dots, \phi_N)$$

$$K_2 = (\omega_1, \omega_2, \dots, \omega_N).$$

Two-point crossover—Two-point crossover randomly selects two locations q_1 and q_2 . First, q_1 is randomly selected such that $q_1 \in \mathbb{Z} : 0 < q_1 < N - 1$. Then, q_2 is randomly selected such that $q_2 \in \mathbb{Z} : q_1 < q_2 < N$. The decision variable vector of the resulting child, K_C , acquires the first $0 \dots q_1$ decision variables from the first parent, K_1 , the middle $q_1 + 1 \dots q_2$ decision variables from the second parent, K_2 , and the last $q_2 + 1 \dots N$ decision variables from the first parent, K_1 . The resulting child is, therefore, defined as

$$K_C = (\phi_0, \dots, \phi_{q_1}, \omega_{q_1+1}, \dots, \omega_{q_2}, \phi_{q_2+1}, \dots, \phi_N).$$

For details on Single-Point and Arithmetic Crossover, alternative crossover operators we tested, see Underwood [56].

Selection and crossover are used to create new individuals to replace every individual in the current population except the most-fit individual. This concept of protecting the best individual in the population is referred to as elitism. Elitism has been shown to substantially improve the performance of GAs [59].

Mutation—After creating these new individuals, the mutation operator is applied to every decision variable in every individual in the population, in order to introduce entropy into the population. The mutation rate is a parameter that governs how much entropy is introduced. The mutation rate, $\mathcal{M}$, is the probability that a single decision variable will be randomly altered. For each decision variable, we generate a random variate, u , such that $u \sim \text{UNIFORM}(0, 1)$. If $u < \mathcal{M}$, then the decision variable is randomly assigned a value in its domain. Application of the mutation operator promotes discovery of new solutions not already encoded in some combination of individuals in the population. It combats the tendency of the heuristic to become trapped in local optima.

In order for the GA to heuristically search a large solution space, it typically must be run for many generations. Each generation requires multiple simulations of individual screening policies in order to estimate expected QALYs. Longer simulations produce lower variance estimates but require more compute time. To balance the need for statistical precision with the need for computational efficiency, we use the Rinott [46] ranking and selection procedure to determine the appropriate sample size for simulating individuals in the population [46]. Rinott's procedure is a two-stage procedure to select the single best of several alternatives (where alternatives with means differing by less than a prescribed "indifference zone" are considered equivalent). Under reasonable assumptions, it provides a statistical guarantee on the probability of correct selection, $P(CS)$. To summarize the procedure we define the following notation: N_0 is the stage-one sample size for variance estimation, and h is a constant computed as

Algorithm 1

GeneticAlgorithm()

```

1  begin
2    Generate initial population of policies;
3    while stopping criteria not met do
4      foreach policy  $K_i$  do
5        Use Rinott method to determine sample size,  $N_p$  for desired  $P(CS)$ ;
6        Simulate policy using determined sample size;
7        Record  $J(K_i)$ ;
8      end
9      BestPolicy  $\leftarrow$  policy with greatest expected QALYs;
10      $J^{\max} \leftarrow J(\text{BestPolicy})$ ;
11     foreach policy  $K_i$  do
12       if  $J(K_i) \approx J^{\max}$  then
13         Parent1  $\leftarrow$  TournamentSelect();
14         Parent2  $\leftarrow$  TournamentSelect();
15         NewPolicy  $\leftarrow$  Crossover(Parent1, Parent2);
```

```

16   NewPolicy ← Mutate(NewPolicy);
17    $D_i \leftarrow$  NewPolicy;
18   end
19   end
20   end
21   return BestPolicy;
21   end

```

the root of an equation (see Rinott [46]) that includes a multi-dimensional integral. The steps of the Rinott procedure are:

1. Simulate each individual K_i using sample size N_0
2. Simulate each individual K_i using sample size $N_i - N_0$, where

$$N_i = \max \left\{ N_0, \left(\frac{h}{\delta^*} \right)^2 S_{K_i}^2 \right\}, \delta^* \text{ defines the}$$

Algorithm 2

TournamentSelect()

Output: Selected policy

```

1   begin
2    $K_A \leftarrow$  randomly selected policy from entire population of policies;
3    $K_B \leftarrow$  randomly selected policy from entire population of policies;
4   if  $J(K_A) > J(K_B)$  then
5       return  $K_A$ ;
6   end
7   else
8       return  $K_B$ ;
9   end
10  end

```

Algorithm 3

Crossover(Parent1, Parent2)

Input: Parent policies which reproduce to form a new policy

Data: L is the length of an individual's string representation (i.e. length of decision variables vector)

Output: New policy formed by reproduction

```

1   begin
2    $cxpoint1 \leftarrow$  random integer  $\in [1, L - 1]$ ;
3    $cxpoint2 \leftarrow$  random integer  $\in [cxpoint1 + 1, L]$ ;
4   for  $i = 1$  to  $cxpoint1$  do
5       NewPolicy[i] ← Parent1[i];
6   end
7   for  $i = cxpoint1 + 1$  to  $cxpoint2$  do

```

```

8   NewPolicy[i] ← Parent2[i];
9   end
10  for  $i = c\text{point}2 + 1$  to  $L$  do
11   NewPolicy[i] ← Parent1[i];
12  end
13  return NewPolicy;
14 end

```

indifference zone, and $S_{K_i}^2$ is the sample variance of the set of N_0 simulations of K_i

3. Compute the expected QALYs for each individual K_j , where $J(K_j)$ is the expected QALYs of K_j
4. Select as best the individual K_j with largest expected QALYs

Pseudocode for the GA is shown in Algorithm 1; the GA's sub-routines are shown in Algorithms 2–4. At each generation of the GA, the Rinott procedure selects the best policy with probability $P(CS)$ from the population. In addition to providing a lower-bound on

Algorithm 4

Mutate(Policy)

Input: Policy to be mutated

Data: L is the length of an individual's string representation (i.e. length of decision variables vector); $\mathcal{M}$ is the mutation rate.

Output: Mutated policy

```

1  begin
2  for  $i = 1$  to  $L$  do
3    $x \leftarrow$  uniform random variate  $\in [0, 1]$ ;
4   if  $x \leq \mathcal{M}$  then
5    Policy[i] ← randomly selected PSA threshold  $\in \{0, 0.5, \dots, 10, \infty\}$ ;
6   end
7  end
8  return Policy;
9  end

```

the $P(CS)$ in the context of selecting the best individual at a given generation, the Rinott procedure also provides a lower bound on the $P(CS)$ in the context of the pairwise comparisons made during tournament selection.

5 Results

In this section, we present computational experiments to evaluate the performance of the GA; then we present results from using the GA to solve the prostate cancer screening problem; and, finally, we compare screening policies from the literature to policies generated by the GA. Our implementation of the simulation optimization method of Section 4 was written in C++ using OpenMP for parallel processing. All numerical experiments were performed on a Linux server with 2 Quad-Core Intel Xeon E5420 2.5 GHz CPUs and 16 GB shared RAM.

5.1 Data

Table 3 displays the values and their sources for the parameters used in our model. These values are based on those used in Zhang et al. [63]. Refer to Tables 1 and 2 for definitions of the parameters. We estimate cancer mortality rates from Surveillance, Epidemiology, and End Results (SEER) data [40], and we base our PSA sampling on a longitudinal data set from Olmsted County, MN. The actual numbers of observations placed into the $\{T \times P\}^{NC}$ and $\{T \times P\}^{Ca}$ bin sets defined in Section 4.2 are listed in Tables 4 and 5, respectively. Each simulated patient started at age 40 years, with the maximum epoch, N , at 100 years.

5.2 Performance of the GA

The various parameters of the GA were set according to performance in preliminary experiments. Tournament selection exhibited better performance than fitness proportionate selection. Two-point crossover exhibited better performance than single-point or arithmetic crossover. A mutation rate of 0.001 performed best, and is recommended in the literature [7]. The GA performed well with population size 30, where the initial population comprises 8 policies from Ross et al. [47] and 22 policies wherein each threshold was uniformly randomly selected from the range of allowable thresholds. The GA appeared to converge after 2000 generations; this convergence was robust to the probability of correct selection, $P(CS)$, in the Rinott procedure.

The evaluation (simulation) step of the GA was implemented in parallel on a shared-memory server with multiple processing cores. The running time of the GA was compared to the number of available cores to estimate the speedup obtained. The running time required by the GA to complete 25 generations using 1, 2, 4, and 8 cores was 1,347, 674, 359, and 178 s, respectively. Thus, the speedup is almost perfectly linear in the number of cores.

We also observed the running time of the GA for different probabilities of correct selection. The running time required by the GA to complete 25 generations using $P(CS)$ equal to 0.2, 0.5, 0.7, and 0.9 was 91, 135, 178, and 251 s, respectively. Thus, $P(CS)$ significantly affects the running time of the GA. As shown in Fig. 5, the quality of the best solution generated by the GA improved from using $P(CS) = 0.2$ to peak at using $P(CS) = 0.7$; thus, using $P(CS) = 0.7$ generated the highest-quality solution most efficiently.

5.3 Base case results

Figure 6 shows the best policy generated by the GA. Because of the distribution of PSA values and the randomness inherent to any GA, there are screening thresholds in Fig. 6 which have little or no effect on the expected QALYs of the shown policy. To separate noise from signal, we systematically change each screening threshold to ∞ (denoted by “inf” on the y-axis of the graphs); thus, we essentially remove each screening threshold. For each threshold, if the expected QALYs of the policy after removing the threshold is not statistically lower ($\alpha = 0.05$) than the expected QALYs of the policy with no thresholds removed (i.e. the policy with noise), then that threshold is removed. This heuristic is performed consecutively on each threshold from age 40 to 100. Figure 7 shows the improved policy generated by the GA after this process. Hereafter, all references to the policy generated by the GA refer to the GA-based policy with noise removed.

As shown in Figs. 8 and 9, the GA generated policy significantly outperforms previously published policies in terms of the QALYs impact on the entire male population and on the most-affected subpopulation (males who ultimately develop cancer between the age of 40 and 50). The previously published policies are taken from Ross et al. [47] and are provided in Table 6. Figure 9 in particular illustrates the fact that relatively small benefits in QALYs for the general population correspond to very significant benefits to the portion of the

population that actually develops prostate cancer. For example, the GA policy results in an increase in 0.81 QALYs over the best of the Ross policies for the most-affected subpopulation.

We examined the effect on the GA-based policy of increasing the instantaneous disutility associated with biopsy. The relative performance of the Ross policies and the GA-based policy were robust to increases in the disutility of biopsy. Thus, the comparatively low thresholds in the GA-based policy are not the result of underestimating the disutility of biopsy.

To put the GA-based policy in perspective, we compared it to simpler uniform threshold policies. Using Fig. 7 as a guide, we explored annual screening policies, screening from age 50 to age 70, over varied uniform thresholds. The 95% confidence intervals around the E[QALYs] of the uniform threshold policies for thresholds 0, 0.5, 1, 1.5, 2, 2.5, 3, 3.5, and 4.0 ng/mL are listed in Table 7. The 95% confidence interval around the E[QALYs] of the policy in Fig. 7 is [38.46, 38.46]. Thus, the uniform policies using thresholds 1.5 or 2.0 ng/mL are close to the GA-based policy.

Despite the fact that the uniform policies perform nearly as well as the GA-based policy, there is merit to using the GA. First, the fact that a simpler uniform policy is nearly optimal can only be asserted with any confidence because we used the GA. The GA allows us to exclude a large number of policies with confidence that we are not inadvertently excluding a significantly better policy. Second, the GA-based policy is in fact better, albeit marginally, than the best of the uniform policies. Practically speaking, the uniform policies may be superior in light of their ease of implementation and the small difference in QALYs. However, we present both alternatives so that the reader may form his or her own opinion.

To more explicitly examine the question of screening frequency, we explore the effect on E[QALYs] of varying the screening frequency of the 1.5 ng/mL uniform threshold policy. The 95% confidence intervals around the E[QALYs] of the 1.5 ng/mL uniform policy for screening intervals 1, 2, 3, 4, and 5 years are listed in Table 8. The E[QALYs] of the evaluated policy decrease as the screening interval exceeds 1 year.

5.4 Sensitivity analysis

We performed one-way sensitivity analysis on the nine parameters from Table 3. For each of the 18 parameter configurations, the GA was run for 1,000 generations, and the best policy produced by the GA after the final generation was then evaluated under the same parameter configuration to obtain tight confidence intervals (using samples of 1×10^8 simulated patients).

Parameter w_p , the prostate cancer (PCa) incidence rate, was varied using the upper and lower bounds from Bubendorf et al. [10] in Table 9. Parameter f , the prostate biopsy detection rate, was varied using 0.64 and 0.96. Parameter z_p , the PCa mortality rate for patients with metastasized PCa, was varied using 0.07 from Messing et al. [39] and 0.37 from Aus et al. [6]. Parameter μ , the instantaneous disutility of prostate biopsy, was varied using 0.01 and 0.1. Parameter ϵ , the decrement of living in the treatment state, was varied using 0.05 and 0.24 from Bremner et al. [9]. Parameter γ , the decrement of living with metastasized PCa, was varied using 0.15 from Krahn et al. [36] and 0.46 from Sandblom et al. [49]. Parameters d_p , e_p , and b_p respectively the death rate from causes other than PCa, the rate of metastasis for patients with undiagnosed PCa, and the rate of metastasis for patients being treated for PCa, were varied using $\pm 20\%$ of the values from Table 3.

The results of the sensitivity analysis are shown in Table 10. These results reveal that the policy is least sensitive to changes in the detection rate of prostate biopsy, f . On the other hand, variation in estimates of the age-dependent mortality rate from causes other than prostate cancer, d_b , can substantially influence the choice of screening policy. Many of the other parameters, including disutilities of metastatic PCa, treatment, and biopsy, have substantial effects when varied within plausible ranges. Of the disutility parameters, those for metastatic prostate cancer, γ , and prostate cancer treatment, ϵ , have the most significant effect. In comparison, the effect of the disutility for biopsy, μ , is much lower.

There were no notable differences in pattern among the policies produced by the GA caused by varying the parameters in the sensitivity analysis, except for the policies produced by varying ϵ and μ . At the high value of ϵ , screening was indicated only between the ages of 50 and 60, whereas at the low value of ϵ , screening was indicated between the ages of 50 and 75 at slightly higher thresholds than in the base case policy. At the high value of μ , thresholds were slightly higher than in the base case policy, whereas at the low value of μ , thresholds were slightly lower than in the base case policy and screening was only indicated between the ages of 50 and 65.

6 Conclusion

The performance results of our GA demonstrate that using a ranking-and-selection procedure, such as the Rinott method, can significantly reduce the running time of the GA. In our GA, setting the $P(CS)$ for the Rinott method to 0.7 reduced the running time to approximately 71% of the running time required when $P(CS)$ was set to 0.9. The quality of the computed policy at both values of $P(CS)$ were equivalent. Hence, our research provides evidence that ranking-and-selection can be an effective tool for balancing the tradeoff between exploration and exploitation in GAs.

Differences in expected QALYs of different PSA screening policies were often less than 0.01. Because of this small difference in QALYs, our GA performed better using tournament selection than using fitness proportionate selection. In fitness proportionate selection, if the fitnesses of the individuals are close, then fitness proportionate selection becomes nearly equivalent to uniform random selection. For example, using fitness proportionate selection, if the fitnesses of all the individuals are equal, then each individual will be selected with equal probability. However, uniform random selection does not provide any guarantee of long-run improvement in fitness, because there is no bias toward more fit individuals. We believe the tournament selection method outperformed the fitness proportionate selection method for precisely this reason. Differences in expected QALYs of medical intervention in other areas are often as small as or smaller than the differences for our problem. For example, Wright and Weinstein [62] estimated the expected perperson benefit of a widely adopted population-based vaccination program to prevent measles and rubella to be 0.008 QALYs. Therefore, our research may guide future disease-treatment studies implementing GAs to examine impact on quality of life.

Analyzing the previously published policies for PSA screening presented in Ross et al. [47] and listed in Table 6, we see from Fig. 8 that policy H outperforms policy G. Policy H is simply a more frequently-screened version of policy G. This may be viewed as evidence that screening less frequently than annually can reduce the expected QALYs of the population. Comparing policy E (common U.S. policy) to policy H, we see little difference in the expected QALYs for the population. The only difference between the two policies is that policy H has two additional screening events, screening once at age 40 and once at age 45. If one interprets this to mean that the screening events at age 40 and at age 45 are insignificant with respect to QALYs, then applying this interpretation to policy B and policy C renders

the only difference between policy B and policy C the frequency of screening. Therefore, the fact that policy C outperforms policy B in Fig. 8 may be taken as evidence, along with the analysis of policy H compared to policy G, that screening less frequently in the frequency range of once per year to once per 5 years reduces the expected QALYs of the population.

Our computed policy suggests screening for typical patients should occur annually starting at age 50 and end at age 70. This suggested starting age for screening is in keeping with many guidelines, however, the ending age for screening is 5 years earlier than many guidelines [47]. It is possible that screening from age 70 to 75 is effective at reducing mortality rates or other measures, but in terms of QALYs, screening beyond age 70 may not be warranted based on the results of our model.

Additionally our results indicate there may be benefits to lowering the PSA thresholds from 4.0 ng/mL to somewhere between 1.5 ng/mL and 3.0 ng/mL. In Fig. 7, the thresholds for the computed policy are on average between 1.5 ng/mL and 3.0 ng/mL. This is similar to, but somewhat more conservative than, the standard U.S. policy of screening annually from age 50 to 75 at threshold 4.0 ng/mL [47]. This computed policy is most similar to policy F from Table 6. Not surprisingly, then, policy F outperforms the rest of the previously published policies in Table 6 with respect to the expected QALYs of the population. Although the policy computed by the GA is variable in the PSA thresholds, which is due in part to the fact that GAs are probabilistic search algorithms, it is clear from Fig. 7 that the computed policy suggests an average threshold substantially lower than the common threshold of 4.0 ng/mL. This may be interpreted as evidence that lowering the threshold for the standard U.S. policy could increase the expected QALYs of the population.

From sensitivity analysis, we observe that the expected QALYs of a policy produced by the GA is most sensitive to variation in the age-dependent rate of death from causes other than prostate cancer, and the prostate cancer incidence rate. The parameter having the least effect is the detection rate of the prostate biopsy procedure.

The policy recommendations made in this paper cannot be guaranteed to be optimal for individual patients. Rather, the recommendations are intended to improve the average QALYs resulting from application of the policy as a population-wide screening procedure for a typical patient. As discussed in Section 3.2, published guidelines have prescribed population-wide screening policies for the early detection of prostate cancer. Our results provide evidence suggesting potential benefits in terms of expected QALYs based on the use of the GA policy. However, the optimal policy is sensitive to model input parameters, such as the disutility of treatment. Thus, our results suggest a need for further investigation of how best to estimate these parameters for individual patients.

There are limitations to our study. First, part of the success of the GA may be due to the fact that the GA is trained on the data set that is subsequently used to develop the simulation model to evaluate each policy. In future work, we intend to evaluate the GA-produced policy on an out-of-sample data set. Second, the only treatment our model incorporated is radical prostatectomy because the data for prostatectomy was readily available. When and if more data on other treatment modalities becomes available, our model could easily be extended to incorporate them. Third, our patient and PSA screening data is derived from a single Caucasian population with arguably better access to healthcare than the average U.S. male has. With data that represents a more diverse population, important insights could be found regarding the impact on high-risk populations (e.g., African American males, males with family history of prostate cancer). Fourth, our screening model prescribes decisions up to the first biopsy, and, thus, our findings do not inform the screening decision process

following an initial negative biopsy. Fifth, because of the limited data available on digital rectal examinations (DREs), the screening actions in our model are solely dependent on PSA observations, whereas in practice a physician will also perform a DRE to detect abnormalities of the prostate gland. Finally, our PSA sampling procedure may not accurately represent PSA levels in some men. Although our longitudinal data set is among the largest currently available, as shown in Tables 4 and 5 we have fewer data points for younger men and for men with prostate cancer. Thus, there are some unavoidable sources of bias in our data. Our model could be extended to address these limitations with the future availability of relevant data, and, therefore, our model forms a foundation for future studies in these directions.

Acknowledgments

This material is based in part upon work supported by the National Science Foundation under Grant Number CMMI 0844511. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. This study was also made possible by the Rochester Epidemiology Project (Grant #R01 - AG034676-47 from the National Institute of Aging). We wish to thank two anonymous reviewers for their constructive comments, which were instrumental in improving this manuscript.

References

1. American Cancer Society. NCCN guidelines for patients: prostate cancer. 2010. <http://www.cancer.org/Cancer/ProstateCancer/MoreInformation/ProstateCancerEarlyDetection/prostate-cancer-early-detection-ac-s-recommendations>
2. American Urological Association. Prostate-specific antigen (PSA) best practice policy. *Oncology*. 2000; 14:277–280.
3. Andriole GL, et al. Mortality results from a randomized prostate-cancer screening trial. *N Engl J Med*. 2009; 360(13):1310–1319. [PubMed: 19297565]
4. Arias E. United states life tables. *Natl Vital Stat Rep*. 2006; 58(21):1–40.
5. Ashley RA, Inman BA, Routh JC, Mynderse LA, Gettman MT, Blute ML. Reassessing the diagnostic yield of saturation biopsy of the prostate. *Eur Urol*. 2008; 53(5):976–981. [PubMed: 17997028]
6. Aus G, Robinson D, Rosell J, Sandblom G, Varenhorst E. Survival in prostate carcinoma-outcomes from a prospective, population-based cohort of 8887 men with up to 15 years of follow-up: results from three countries in the population-based national prostate cancer registry in Sweden. *Cancer*. 2006; 103(5):943–951. [PubMed: 15651057]
7. Baluja, S.; Caruana, R. Removing the genetics from the standard genetic algorithm. In: Kuhl, ME.; Steiger, NM.; Armstrong, FB.; Joines, JA., editors. *Proceedings of the twelfth international conference on machine learning*; 1995.
8. Bastian PJ, Mangold LA, Epstein JI, Partin AW. Characteristics of insignificant clinical T1c prostate tumors. A contemporary analysis. *Cancer*. 2004; 101(9):2001–2005. [PubMed: 15372478]
9. Bremner KE, Chong CAKY, Tomlinson G, Alibhai SMH, Krahn MD. A review and meta-analysis of prostate cancer utilities. *Med Decis Mak*. 2007; 27:288–298.
10. Bubendorf L, Schöpfer A, Wagner U, Sauter G, Moch H, Willi N, Gasser TC, Mihatsch MJ. Metastatic patterns of prostate cancer: an autopsy study of 1589 patients. *Human Pathol*. 2000; 31(5):578–583. [PubMed: 10836297]
11. Buchholz, P.; Thümmel, A. Enhancing evolutionary algorithms with statistical selection procedures for simulation optimization. In: Kuhl, ME.; Steiger, NM.; Armstrong, FB.; Joines, JA., editors. *Proceedings of the 2005 winter simulation conference*; 2005. p. 842–852.
12. Chen C-H. A lower bound for the correct subset-selection probability and its application to discrete-event system simulations. *IEEE Trans Automat Contr*. 1996; 41(8):1227–1231.
13. Chen EJ, Lee LH. A multi-objective selection procedure of determining a pareto set. *Comput Oper Res*. 2009; 36:1872–1879.

14. Chhatwal J, Alagoz O, Burnside ES. Optimal breast biopsy decision-making based on mammographic features and demographic factors. *Oper Res.* 2010; 58(6):1577–1591. [PubMed: 21415931]
15. Chon CH, Lai FC, McNeal JE, Presti JC. Use of extended systematic sampling in patients with a prior negative prostate needle biopsy. *J Urol.* 2002; 167(6):2457–2460. [PubMed: 11992057]
16. Deb, K.; Agrawal, S.; Pratap, A.; Meyarivan, T. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II. In: Schoenauer, M.; Deb, K.; Rudolph, G.; Yao, X.; Lutton, E.; Merelo, J.; Schwefel, H-P., editors. *Parallel problem solving from nature PPSN VI. Lecture notes in computer science.* Vol. vol 1917. Berlin: Springer; 2000. p. 849-858. ISBN 978-3-540-41056-0
17. Djavan B, Margreiter M. Biopsy standards for detection of prostate cancer. *World J Urol.* 2007; 25:11–17. [PubMed: 17342490]
18. Djulbegovic M, Beyth RJ, Neuberger MM, Stoffs TL, Vieweg J, Djulbegovic B, Dahm P. Screening for prostate cancer: systematic review and meta-analysis of randomised controlled trials. *BMJ.* 2010; 341:9.
19. Draisma G, Boer R, Otto SJ, van der Cruijsen IW, Damhuis RAM, Schröder FH, de Koning HJ. Lead times and overdiagnosis due to prostate-specific antigen screening: estimates from the european randomized study of screening for prostate cancer. *J Natl Cancer Inst.* 2003; 95(12):868–878. [PubMed: 12813170]
20. Dudewicz EJ, Dalal SR. Allocation of observations in ranking and selection with unequal variances. *Indian J Stat.* 1975; 37(1):28–78.
21. Epstein JI, Sanderson H, Carter HB, Scharfstein DO. Utility of saturation biopsy to predict insignificant cancer at radical prostatectomy. *Urology.* 2005; 66(2):356–360. [PubMed: 16040085]
22. Eskandari, H.; Rabelo, L.; Mollaghasemi, M. Multiobjective simulation optimization using an enhanced genetic algorithm; Proceedings of the 37th conference on winter simulation. Winter simulation conference; 2005. p. 833-841.
23. Etzioni RB, Cha R, Cowen ME. Serial prostate specific antigen screening for prostate cancer: a computer model evaluates competing strategies. *J Urol.* 1999; 162:741–748. [PubMed: 10458357]
24. Etzioni RB, Gulati R, Falcon S, Penson DF. Impact of PSA screening on the incidence of advanced stage prostate cancer in the united states: a surveillance modeling approach. *Med Decis Mak.* 2008; 28(3):323–331.
25. Ferrini R, Woolf SH. American college of preventive medicine practice policy: screening for prostate cancer in American men. *Am J Prev Med.* 1998; 15(1):81–84. [PubMed: 9651646]
26. Fu MC. Optimization for simulation: theory vs. practice. *INFORMS J Comput.* 2002; 14(3):192–215.
27. Ghani KR, Grigor K, Tulloch DN, Bollina PR, McNeill SA. Trends in reporting Gleason score 1991 to 2001: changes in the pathologist's practice. *Eur Urol.* 2005; 47(2):196–201. [PubMed: 15661414]
28. Gold MR, Stevenson D, Fryback DG. HALYs and QALYs and DALYs, oh my: similarities and differences in summary measures of population health. *Annu Rev Public Health.* 2002; 23:115–134. [PubMed: 11910057]
29. Goldsman, D.; Nelson, BJ. Ranking, selection and multiple comparisons in computer simulation; Proceedings of the 1994 winter simulation conference. Winter simulation conference; 1994.
30. Gustafsson L, Adami HO. Optimization of cervical cancer screening. *Cancer Causes Control.* 2008; 3(2):125–136. [PubMed: 1562702]
31. Haas GP, Delongchamps NB, Jones RF, Chandan V, Serio AM, Vickers AJ, Jumbelic M, Threutte G, Korets R, Lilja H, de la Roza G. Needle biopsies on autopsy prostates: sensitivity of cancer detection based on true prevalence. *J Natl Cancer Inst.* 2007; 99(19):1484–1489. [PubMed: 17895474]
32. Inoue K, Chick SE, Chen C-H. An empirical evaluation of several methods to select the best system. *ACM Trans Model Comput Simul.* 1999; 9(4):381–407.
33. Jemal A, Siegel R, Ward E, et al. Cancer statistics, 2009. *CA Cancer J Clin.* 2009; 59(4):225–249. [PubMed: 19474385]

34. Klein T, Palisaar RJ, Holz A, Brock M, Noldus J, Hinkel A. The impact of prostate biopsy and periprostatic nerve block on erectile and voiding function: a prospective study. *J Urol*. 2010; 184(4):1447–1452. [PubMed: 20727540]
35. Koenig LW, Law AM. A procedure for selecting a subset of size m containing the l best of k independent normal populations, with applications to simulation. *Commun Stat Theory Methods*. 1985; 14(3):719–734.
36. Krahn M, Ritvo P, Irvine J, Tomlinson G, Bremner KE, Bezjak A, Trachtenberg J, Naglie G. Patient and community preferences for outcomes in prostate cancer: implications for clinical policy. *Med Care*. 2003; 41(1):153–164. [PubMed: 12544552]
37. Maillart LM, Ivy JS, Ransom S, Diehl K. Assessing dynamic breast cancer screening policies. *Oper Res*. 2008; 56(6):1411–1427.
38. Master VA, Chi T, Simko JP, Weinberg V, Carroll PR. The independent impact of extended pattern biopsy on prostate cancer stage migration. *J Urol*. 2005; 174(5):1789–1793. discussion 1793. [PubMed: 16217288]
39. Messing EM, Manola JY, Kiernan M, Crawford GW, di'Sant Agnese PA, Trump D. Immediate versus deferred androgen deprivation treatment in patients with node-positive prostate cancer after radical prostatectomy and pelvic lymphadenectomy. *Lancet Oncol*. 2006; 7(5):472–479. [PubMed: 16750497]
40. National Cancer Institute. Surveillance epidemiology and end results. 2008. <http://seer.cancer.gov/>
41. National Cancer Institute. Cancer intervention and surveillance modeling network. 2011. <http://cisnet.cancer.gov/prostate/>
42. National Comprehensive Cancer Network. American Cancer Society recommendations for prostate cancer early detection. 2011. <http://www.nccn.com/images/patient-guidelines/pdf/prostate.pdf>
43. Presti JC. Prostate biopsy: how many cores are enough. *Urol Oncol*. 2003; 21:135–140. [PubMed: 12856642]
44. Rabets JC, Jones JS, Patel A, Zippe CD. Prostate cancer detection with office based saturation biopsy in a repeat biopsy population. *J Urol*. 2004; 172(1):94–97. [PubMed: 15201745]
45. Ransohoff DF, Collins MM, Fowler FJ Jr. Why is prostate cancer screening so common when the evidence is so uncertain? A system without negative feedback. *Am J Med*. 2002; 113(8):663–667. [PubMed: 12505117]
46. Rinott Y. On two-stage selection procedures and related probability-inequalities. *Commun Stat Theory Methods*. 1978; A7(8):799–811.
47. Ross KS, Carter HB, Pearson JD, Guess HA. Comparative efficiency of prostate-specific antigen screening strategies for prostate cancer detection. *J Am Med Assoc*. 2000; 284(11):1399–1405.
48. Sanda MG, Kaplan ID. A 64-year-old man with low-risk prostate cancer: review of prostate cancer treatment. *J Am Med Assoc*. 2009; 301(20):2141–2151.
49. Sandblom G, Carlsson P, Sennfl K, Varenhorst E. A population-based study of pain and quality of life during the year before death in men with prostate cancer. *Br J Cancer*. 2004; 90(6):1163–1168. [PubMed: 15026796]
50. Scardino PT, Beck JR, Miles BJ. Conservative management of prostate cancer. *N Engl J Med*. 1994; 330(25):1831. author reply 1831–1831; author reply 1832. [PubMed: 8190171]
51. Schröder FH, WildHagen MF. Screening for prostate cancer: evidence and perspectives. *Br J Urol Int*. 2001; 88:811–817.
52. Schröder FH, et al. Screening and prostate-cancer mortality in a randomized European study. *N Engl J Med*. 2009; 360(13):1320–1328. [PubMed: 19297566]
53. Smith RA, Cokkinides V, Brooks D, Saslow D, Brawley OW. Cancer screening in the united states, 2010: a review of current American Cancer Society guidelines and issues in cancer screening. *CA Cancer J Clin*. 2010; 60(2):99–119. [PubMed: 20228384]
54. Swisher JR, Jacobson SH, Insead EY. Discrete-event simulation optimization using ranking, selection, and multiple comparison procedures: a survey. *ACM Trans Model Comput Simul*. 2003; 13(2):134–154.
55. Thompson IM, Ankerst DP, Chi C, Goodman PJ, Tangen CM, Lucia MS, Feng Z, Parnes HL, Coltman CA. Assessing prostate cancer risk: results from the prostate cancer prevention trial. *J Natl Cancer Inst*. 2006; 98(8):529–534. [PubMed: 16622122]

56. Underwood, DJ. Simulation optimization of prostate cancer screening using a parallel genetic algorithm. 2010. <http://www.lib.ncsu.edu/resolver/1840.16/6412>
57. U.S. Preventive Services Task Force. Screening for prostate cancer: U.S. Preventive Services Task Force recommendation statement. *Ann Intern Med*. 2008; 149(3):185–191. [PubMed: 18678845]
58. U.S. Preventive Services Task Force. Screening for prostate cancer: draft recommendation statement. AHRQ Publication No. 12-05160-EF-2. 2011. <http://www.uspreventiveservicestaskforce.org/uspstf12/prostate/draftrec3.htm>
59. Vasconcelos JA, Ramírez JA, Takahashi RHC, Saldanha RR. Improvements in genetic algorithms. *IEEE Trans Magn*. 2001; 37(5):3414–3417.
60. Welch HG, Albertsen PC. Prostate cancer diagnosis and treatment after the introduction of prostate-specific antigen screening: 1986–2005. *J Natl Cancer Inst*. 2009; 101(19):1325–1329. [PubMed: 19720969]
61. Wolf AMD, Wender RC, Etzioni RB, Thompson IM, D'Amico AV, Volk RJ, Brooks DD, Dash C, Guessous I, Andrews K, DeSantis C, Smith RA. American Cancer Society guideline for the early detection of prostate cancer: update 2010. *CA Cancer J Clin*. 2010; 60(2):70–98. [PubMed: 20200110]
62. Wright JC, Weinstein MC. Gains in life expectancy from medical interventions: standardizing data on outcomes. *N Engl J Med*. 1998; 339(6):380–386. [PubMed: 9691106]
63. Zhang J, Denton BT, Balasubramanian H, Shah ND, Inman BA. Optimization of PSA screening policies: a comparison of the patient and societal perspectives. *Med Decis Mak*. 2011 Prepublished on 20 September 2011.
64. Zhang, J.; Denton, BT.; Balasubramanian, H.; Shah, ND.; Inman, BA. Optimization of prostate biopsy referral decisions. Technical report. Working paper. 2011.

Fig. 1.
Receiver Operating Characteristic (ROC) curve illustrating the imperfect nature of PSA testing based on longitudinal data for a regional population in Rochester, MN

Fig. 2. Health-state transitions diagram where *NC* is no cancer, *C* is cancer not detected, *M* is metastasis, and *D* is death. The treatment state, *T*, is an aggregate state

Fig. 3.
Histogram of PSA observations from Olmsted County, MN data set

Fig. 4.
Flowchart describing the major steps of the simulation optimization method

Expected value of the best policy generated by the GA under different probabilities of correct selection (95% CIs, $n=1 \times 10^8$)

Fig. 5. Performance of the GA using different probabilities of correct selection, $P(CS)$, in the Rinott method

Fig. 6.
Improved policy generated by the GA in raw form (with noise)

Fig. 7.
Improved policy generated by the GA with noise filtered using the heuristic

Fig. 8.
Comparing the Ross policies to the best policy generated by the GA

Fig. 9.
Comparing the effect of the Ross policies and the best GA policy on subpopulation of males who develop cancer between the age of 40 and 50

Table 1

Definitions of reward parameters used in evaluating proposed screening policies

Parameter	Definition
ϵ	Annual utility decrement of living in state T
γ	Annual utility decrement of living in state M
μ	One-time utility decrement associated with prostate biopsy
f	Biopsy detection rate for patients with prostate cancer

Table 2

Definitions of parameters for transition probability matrices

Parameter	Definition
w_t	Annual prostate cancer incidence rate
d_t	Annual death rate from other causes
b_t	Annual metastasis rate for patients in state T
e_t	Annual metastasis rate for patients in state C
z_t	Annual prostate cancer death rate excluding death from other causes for patients in state M
f	Biopsy detection rate for patients with prostate cancer

Table 3

Values of parameters used in our base case analysis and their sources

Parameter	Value	Source
w_t	Age specific	Bubendorf et al. [10]
d_t	Age specific	Arias [4]; National Cancer Institute [40]
b_t	0.006	MCRPR
e_t	0.069	Ghani et al. [27]; Scardino et al. [50]
z_t	Age specific	National Cancer Institute [40]
f	0.8	Haas et al. [31]
ϵ	0.145	Bremner et al. [9]
γ	0.24	Bremner et al. [9]
μ	0.05	Chhatwal et al. [14]

MCRPR is the Mayo Clinic Radical Prostatectomy Registry

Table 4Number of observations in sampling bin sets $\{\mathbb{T} \times \mathbb{P}\}^{NC}$

	[0, 49)	[50, 59)	[60, 69)	[70, ∞)
[0, 1)	9,925	30,414	28,647	27,370
[1, 2)	832	2,673	3,299	3,142
[2, 3)	160	738	1,450	1,763
[3, 4)	72	292	641	1,150
[4, 7)	88	246	570	1,509
[7, 10)	28	56	153	478
[10, ∞)	50	54	115	366

Table 5Number of observations in sampling bin sets $\{\mathbb{T} \times \mathbb{P}\}^{Ca}$

	[0, 49)	[50, 59)	[60, 69)	[70, ∞)
[0, 1)	164	1,077	1,815	2,140
[1, 2)	34	176	223	165
[2, 3)	9	107	180	144
[3, 4)	2	66	177	134
[4, 7)	8	57	218	349
[7, 10)	2	3	42	136
[10, ∞)	0	10	20	126

Table 6

Previously published prostate-cancer screening policies taken from Ross et al. [47]

Label	Policy
A	No screening
B	40–75 (every 5 years): 4.0
C	50–75 (every 2 years): 4.0
D	50–59: 3.5; 60–69: 4.5; 70–75: 6.5
E	50–75: 4.0
F	50–75: 2.5
G	40, 45, 50–75 (every 2 years): 4.0
H	40, 45, 50–75: 4.0

Descriptions follow this format: age(s) in years for PSA testing, colon, threshold in ng/mL; multiple descriptions separated by semicolon. Screening frequency over age intervals is annual unless otherwise stated. Labels in this table are referred to only in graphs

Table 7

95% confidence intervals around the E[QALYs] for annual screening from age 50 to 70 at uniform thresholds

Thresholds	95% CI
0	[38.41, 38.42]
0.5	[38.42, 38.43]
1	[38.44, 38.45]
1.5	[38.45, 38.46]
2	[38.45, 38.46]
2.5	[38.45, 38.45]
3	[38.44, 38.45]
3.5	[38.43, 38.44]
4	[38.42, 38.42]

Table 8

95% confidence intervals around the E[QALYs] for screening from age 50 to 70 at different intervals using threshold 1.5 ng/mL

Intervals (yrs.)	95% CI
1	[38.45, 38.46]
2	[38.44, 38.44]
3	[38.43, 38.43]
4	[38.42, 38.43]
5	[38.41, 38.42]

Table 9Upper and lower bounds on w_t for sensitivity analysis

t	LB on w_t	UB on w_t
40–49	2.00×10^{-4}	5.01×10^{-3}
50–59	1.51×10^{-3}	4.91×10^{-3}
60–69	2.43×10^{-3}	8.52×10^{-3}
70–79	5.22×10^{-3}	1.51×10^{-2}
80	7.12×10^{-3}	1.10×10^{-2}

Table 10

One-way sensitivity analysis on best GA policy for parameters d_t , w_t , γ , e_t , ϵ , z_t , μ , b_t and f

Parameter	Low	High	Gap
d_t	36.664	40.657	3.993
w_t	38.042	38.454	0.412
γ	38.307	38.514	0.207
e_t	38.375	38.547	0.172
ϵ	38.387	38.535	0.148
z_t	38.406	38.510	0.104
μ	38.427	38.483	0.056
b_t	38.447	38.460	0.013
f	38.451	38.453	0.002