

ABSTRACT

UNDERWOOD, DANIEL J. Simulation Optimization of Prostate Cancer Screening Using a Parallel Genetic Algorithm. (Under the direction of Dr. Brian T. Denton.)

We develop a parallel simulation-optimization method to find near optimal PSA screening policies based on the mean expected quality adjusted life years (QALYs) of a large population. In contrast to mortality-based indicators such as life expectancy, QALYs are informed by treatment side-effects, chronic disease, and other non-fatal degradations to quality of life. We seek an optimal balance of the benefits of early detection of prostate cancer and the reduction in QALYs associated with overdiagnosis.

A screening policy's impact on QALYs is examined by using a non-stationary Markov model to simulate a patient's progression through health-states, PSA levels, and treatment. The simulation model is embedded within a parallel genetic algorithm (GA) to search the space of possible PSA screening policies for unrealized optimal policies. A ranking and selection scheme is employed by the GA in order to manage the computational budget. The ranking and selection scheme is designed to balance the trade-off between exploring the large policy space and achieving statistical confidence in the simulation output that is necessary to meaningfully distinguish different screening policies.

Near optimal policies from our model are compared to previously proposed screening policies. The results of sensitivity analyses provide insights into the effects of model parameters. Finally, the best policy from our simulation optimization model is compared to the theoretically best policy from a previously reported POMDP model. The latter results provide an estimate of the benefit or loss resulting from the use of easier-to-implement PSA-threshold policies.

© Copyright 2010 by Daniel J. Underwood

All Rights Reserved

Simulation Optimization of Prostate Cancer Screening
Using a Parallel Genetic Algorithm

by
Daniel J. Underwood

A thesis submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Master of Science

Operations Research

Raleigh, North Carolina

2010

APPROVED BY:

Dr. James R. Wilson

Dr. John W. Baugh

Dr. Brian T. Denton
Chair of Advisory Committee

DEDICATION

To my parents.

BIOGRAPHY

Daniel J. Underwood was born on December 22, 1983, in Raleigh, NC to Dr. Ronald and Mrs. Pam Underwood; the former holds a PhD in Mechanical Engineering from NC State University (NCSU). After receiving a BSc in Chemistry from NCSU in 2006, Daniel began the MSc program in Operations Research at NCSU. Daniel has accepted a teaching assistantship with the Edward P. Fitts Department of Industrial and Systems Engineering at NCSU, and will begin his PhD studies in Industrial Engineering in the Fall of 2010.

ACKNOWLEDGEMENTS

I would like to thank my advisors for their tireless assistance during all stages of my research. I would also like to thank my colleague Jingyu Zhang for his invaluable prior work, Nilay Shah for assembling and providing original data, and Brant Inman for giving expert medical feedback.

TABLE OF CONTENTS

List of Tables	vii
List of Figures	viii
Chapter 1 Introduction	1
Chapter 2 Background on Prostate Cancer Screening	4
Chapter 3 Literature Review	7
3.1 Prostate Cancer Modeling	7
3.2 Simulation Optimization	8
3.2.1 Ranking and Selection	8
3.2.2 Genetic Algorithms	9
3.2.3 Parallel Computing	10
3.3 Contribution to the Literature	10
Chapter 4 Simulation Model	12
4.1 PSA Screening	14
4.2 Accumulation of Reward	15
4.3 Core Health-State Transition Probabilities	15
4.4 Computing Platforms	17
4.5 Numerical Results	18
4.5.1 Data and Model Parameters	18
4.5.2 Discounting QALYs	18
4.5.3 Evaluation of Proposed Screening Policies	19
4.6 Conclusions	22
Chapter 5 Optimization of PSA-Threshold Screening Policies	23
5.1 Optimization Model	24
5.2 Genetic Algorithm	25
5.3 Results	32
5.3.1 Experimental Evaluation of Genetic Algorithm	32
5.3.2 Experimental Evolution of Genetic Algorithm	35
5.3.3 Optimal PSA Screening Policies	35
5.3.4 Clinically Relevant Screening Policies	37
5.3.5 Upper Bounds on the Optimal Policy	40
5.3.6 Multiple Biopsy	40
5.3.7 Sensitivity Analysis	44
5.4 Conclusions	46

Chapter 6 Conclusion	48
6.1 Summary of Conclusions	48
6.2 Implications for Research and Practice	49
6.3 Potential Future Research	50
References	51

LIST OF TABLES

Table 4.1	Parameters for transition probability matrices	16
Table 4.2	Reward parameters used in evaluating proposed screening policies	18
Table 4.3	Transition probabilities used in evaluating proposed screening policies	19
Table 5.1	Precision and running time analysis	32
Table 5.2	Comparison between the best GA policy and the optimal POMDP policy [47]	40
Table 5.3	Sensitivity analysis on parameters of reward function in terms of the effect on the entire U.S. male population	44
Table 5.4	Sensitivity analysis on parameters of reward function in terms of the effect on the subpopulation of U.S males who ultimately develop cancer	45
Table 5.5	Recommendation for PSA-threshold screening policy	46
Table 6.1	Comparison of the expected QALYs for the Ross policies	48
Table 6.2	Recommendation for PSA-threshold screening policy	49

LIST OF FIGURES

Figure 2.1	Receiver Operating Characteristic (ROC) curve illustrating the imperfect nature of PSA testing based on longitudinal data for a regional population in Rochester, MN	6
Figure 4.1	Transitions between health-states	13
Figure 4.2	States considered following treatment	13
Figure 4.3	Histogram of PSA observations from empirical data	14
Figure 4.4	Sensitivity analysis on parameters of reward function for Ross policies	19
Figure 4.5	Ross results comparison for entire U.S. male population	20
Figure 4.6	Ross results comparison for subpopulation of U.S. males who ultimately develop prostate cancer	21
Figure 4.7	Comparison between the Ross policies' effect on the entire U.S. male population and the subpopulation of U.S. males who ultimately develop prostate cancer	21
Figure 5.1	High-level model of GA	25
Figure 5.2	Example roulette wheel illustration	27
Figure 5.3	Single-point crossover illustration	29
Figure 5.4	Two-point crossover illustration	29
Figure 5.5	Convergence of mean as sample size increases	33
Figure 5.6	Comparison of crossover methods	34
Figure 5.7	Comparison of selection methods	34
Figure 5.8	Comparison of mutation rates	35
Figure 5.9	Evolution of GA showing period of sharpest incline in fitness improvement from generation to generation	36
Figure 5.10	Evolution of GA showing convergence to optimum	36
Figure 5.11	Optimal PSA screening policy discovered by GA in raw form	37
Figure 5.12	Optimal PSA screening policy discovered by GA in simplified form	38
Figure 5.13	Comparing the Ross policies to the optimal GA policy (both the raw version and the simplified version) in terms of the effect on the entire U.S. male population	39
Figure 5.14	Comparing the Ross policies to the optimal GA policy in terms of the effect on the subpopulation of U.S. males who ultimately develop prostate cancer	39
Figure 5.15	Optimal PSA screening policy under maximum of 3 biopsies per patient discovered by GA in raw form	41
Figure 5.16	Optimal PSA screening policy under maximum of 3 biopsies per patient discovered by GA in simplified form	41
Figure 5.17	Comparing the Ross policies, under maximum of 3 biopsies per patient, to the optimal GA policy (both the raw version and the simplified version) in terms of the effect on the entire U.S. male population	42

Figure 5.18 Comparing the Ross policies, under maximum of 3 biopsies per patient, to the optimal GA policy in terms of the effect on the subpopulation of U.S. males who ultimately develop prostate cancer 43

Prostate cancer is one of the most common cancers affecting males. Statistics show for males in the United States during 2009 that prostate cancer was responsible for 25% of new cancers and 9% of deaths from cancer (excluding skin cancers) [28]. Clinicians screen for prostate cancer by monitoring the serum-level of prostate-specific antigen (PSA) in males. PSA is a protein produced by cells in the prostate gland; PSA is also a continuous biomarker that is correlated with prostate cancer. Males typically have low levels of PSA in their blood. High PSA levels often point to the existence of cancerous cells in the prostate. However, PSA screening is an imperfect test, resulting in false positives, because increased PSA levels can result from benign prostatic hyperplasia (BPH), also known as benign prostate enlargement. Furthermore, because the all-cause mortality rate increases with age, and because prostate cancer is a slow growing cancer, PSA screening loses its utility in older patients. The age at which to discontinue screening is central to the debate on PSA screening.

When a physician suspects a patient has prostate cancer due to high PSA levels in the blood (usually accompanied by confirming results from a digital rectal examination), the patient is referred for biopsy. During a prostate biopsy, long needles are inserted into the prostate from diverse angles in order to extract 12 cores (samples) for later inspection under a microscope. The insertion of the needles is guided by a transrectal ultrasound, which requires a device to be inserted into the rectum so that the physician can image the prostate area and accurately place the needles into the prostate gland. Patients generally view this process as invasive or discomforting. Furthermore, prostate biopsy is an imperfect test for determining the presence of prostate cancer; Haas et al. [22] estimate that the prostate cancer detection rate for prostate

biopsy is 0.8.

Because the treatment options following a positive PSA test can cause serious negative side-effects [40], reducing the likelihood of false positives is naturally of great interest to researchers, patients, and practitioners.

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 rate of death from prostate cancer [4]. Results from The European Randomized Study of Screening for Prostate Cancer, on the other hand, indicate that PSA screening *does* reduce the rate of death from prostate cancer [41]. The heterogeneity of results from controlled trials on the impact of PSA screening makes this subject particularly suitable for simulation analysis.

The problem of determining the best PSA screening policy is controversial and difficult to solve. Physicians typically consider a discrete set of clinically relevant ranges of PSA. Assuming *arguendo* that there are a total of 5 different possible PSA threshold levels (in reality PSA is a continuous measure), that screening may take place no more frequently than every 3 months, and that screening may begin at age 40, then there are $5^{160} = 6.8 \times 10^{111}$ possible screening policies. Obviously it is not possible to conduct randomized controlled trials to evaluate and compare all possible screening policies. Even using a computer to systematically simulate and evaluate all the possible screening policies is not feasible (e.g. if a computer can evaluate 1,000,000 screening policies per second, it would take 2.2×10^{98} years to evaluate all policies).

The problem of determining the best PSA screening policy demands sophisticated optimization techniques. This optimization problem is difficult for multiple reasons. Because the problem space is stochastic and non-convex, methods such as stochastic linear programming and dynamic programming are not suitable. Due to limited data, the probability distribution of PSA levels is difficult to fit analytically and, therefore, the underlying distribution must be discretized, which substantially obfuscates our understanding of the true problem structure. Finally, based on results of randomized controlled trials, referenced above, expected differences among screening policies are small, requiring large numbers of samples to achieve sufficiently tight confidence intervals.

We search for near-optimal PSA screening policies by exploring the effect of factors such as PSA threshold levels, screening intervals, and treatment decisions on a patient's expected quality adjusted life years (QALYs). QALYs are a means of assigning value to a range of health outcomes from death to perfect health; in contrast to mortality-based indicators such as life expectancy, QALYs are informed by treatment side-effects, chronic disease, and other non-fatal degradations to quality of life [19].

We examine a given policy's impact on QALYs by using a non-stationary Markov model to

simulate a patient's progression through core health-states, PSA levels, and treatment. Next, the simulation model is embedded within a parallel genetic algorithm (GA) to search the space of possible PSA screening policies for optimal policies. A well known ranking and selection scheme is adapted by the GA in order to manage the computational budget. The ranking scheme operates to balance the competing interests in exploring the problem space and exploiting the statistical precision associated with a given set of screening policies. Finally, a simplification procedure is used to change the best policy from the GA into a nearly optimal policy that is more clinically practicable.

The remainder of this thesis is structured as follows: Chapter 2 gives background on prostate cancer screening; Chapter 3 reviews the relevant literature; Chapter 4 presents the simulation model; Chapter 5 presents ranking-and-selection as a tool for computational efficiency, and then describes the genetic algorithm which uses the simulation model to evaluate solutions; and Chapter 6 recapitulates the results from prior chapters and presents conclusions inferred from those results.

Background on Prostate Cancer Screening

Screening for prostate cancer is done primarily by PSA screening. The goal of PSA screening is to detect the cancer while it is localised to the prostate gland so that treatment may be applied to prevent the cancer from spreading to other organs. There are many different screening policies in practice across the world. Each policy defines (or suggests) a frequency for screening and the threshold level of PSA that, when reached, constitutes a positive test for prostate cancer.

Authorities in the United States are not unanimous regarding PSA screening policy recommendations. The American Urological Association and The American Cancer Society recommended annual PSA tests for men above age 50, and starting even earlier for men in high risk categories [4]. The National Comprehensive Cancer Network recommended an algorithm to determine screening policy based on various risk factors [4], such as family history of prostate cancer. The American College of Physicians and the American College of Preventative Medicine recommended *against* routine PSA screening. The U.S. Preventative Services Task Force recommended that males should only be screened under the age of 75 [4]. 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.

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 threshold 3.0 ng/mL. However, Finland, Italy, Holland, and Belgium use threshold 4.0 ng/mL, like the United States. Most European countries use 4-year screening intervals, although Sweden uses 2-year intervals.

The prostate biopsy is the gold standard for detecting prostate cancer. The current standard

method of biopsy is the transrectal ultrasound (TRUS)-guided prostate biopsy. Using an ultrasonographic device inserted into the rectum, the physician performing the biopsy is able to produce an imaging of the pelvic region. Using this imaging as a guide, needles are introduced into the prostate from different vantage points to extract cores (samples). Physicians typically extract 12 cores for analysis [37]. Each core is subsequently analyzed and characterized on the basis of any cancerous mass detected. For a detailed discussion of prostate biopsy standards, see Djavan and Margreiter [14].

Sanda and Kaplan [40] provide an extensive case-based review of prostate cancer treatment, a summary of which we introduce here. The most common treatment modalities are prostatectomy, radiotherapy, and brachytherapy. Prostatectomy is a surgical procedure where a portion or all of the prostate gland is removed. Radiotherapy requires a beam of radiation to be directed from an external source into the prostate gland, and brachytherapy requires a radioactive substance to be implanted into and beside the cancerous regions. Physicians also recognize *active surveillance* as an effective treatment alternative for prostate cancer cases with low risk of metastasis. Active surveillance is essentially deferred primary treatment coupled with more aggressive monitoring. For intermediate- and high-risk prostate cancers, the use of radiotherapy is often combined with adjuvant androgen suppressive therapy. Brachytherapy is limited to low- and intermediate-risk prostate cancers. Prostatectomy is a viable treatment option for many patients with prostate cancer. In all cases, the method of treatment should be chosen after a careful discussion between patient and physician about the risks associated with each treatment and about the patient's personal preferences.

Medical tests for diseases have different levels of *sensitivity* and *specificity*. Sensitivity quantifies how reliable the test is at detecting the disease, whereas specificity quantifies how accurate the test is at identifying people without the disease [31]. More formally, sensitivity defines the probability of a positive test result given that the disease is actually present. Specificity defines the probability of a negative test result given that the disease is not in fact present. The results from a PSA test are not easy to interpret. Rather than having simple positive or negative outcomes, like tests for pregnancy, PSA tests report values which must be interpreted in light of specified threshold values. These sensitivity and specificity quantities are compared by use of Receiver Operating Characteristic (ROC) curves. Physicians use ROC curves to evaluate the performance of diagnostic tests [36]. The ROC curve for PSA testing, from Zhang et al. [46], is shown in Figure 2. The origin of ROC analysis comes from the early 1950s when a group of mathematical psychologists developed the means to quantify the ability of a radar to detect approaching aircraft [34]. Since then, ROC analysis has been applied to many fields.




Figure 2.1: Receiver Operating Characteristic (ROC) curve illustrating the imperfect nature of PSA testing based on longitudinal data for a regional population in Rochester, MN

This review surveys the major solution methodologies we employ to optimize PSA screening policies. We discuss prostate cancer modeling, and simulation in the field of medical research and disease modeling. We then discuss elements needed for a simulation-optimization approach to optimize PSA-based screening for prostate cancer. Finally, we provide a brief review of parallel programming concepts that are relevant to this thesis.

3.1 Prostate Cancer Modeling

A group of 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.¹ Using results from the Rotterdam section of the European Randomized Study of Screening for Prostate Cancer, Draisma et al. [15] modeled the development of prostate cancer up to the point of diagnosis as a Markov process with states defined by cancer stage and cancer differentiation grade, and produced evidence that PSA screening at intervals larger than 1 year is effective.

Ross et al. [39] compared various PSA screening policies using a Monte-Carlo simulation based on a Markov model of the natural progression of prostate cancer, ultimately concluding that the standard U.S. policy is not optimal on the basis of prostate cancer mortality rates. Zhang et al. [46] searched for optimal biopsy referral policies by solving a partially-observable Markov decision process using a fixed-finite-grid approximation with the goal of maximizing

¹See <http://cisnet.cancer.gov/prostate/>

expected quality-adjusted life years.

Researchers have also studied other forms of cancer. For instance, Maillart et al. [32] formulated a partially-observable Markov decision process to evaluate screening policies for the detection of breast cancer—a problem sharing many features with the prostate cancer problem—and found an upper-bound on the optimal age to begin screening. Similarly, Chhatwal et al. [11] formulated the problem of when to biopsy for breast cancer as a finite-horizon discrete-time Markov decision process and concluded that the decision to biopsy is age dependent.

3.2 Simulation Optimization

Simulation optimization is any instance of optimization where the objective function is evaluated by simulating a stochastic system; the simulator is effectively a so-called *black box* that returns a value while there is no clear analytic expression of the function. Stochastic problems which lack properties such as convexity and optimal substructure are not amenable to optimization methods such as stochastic programming or stochastic dynamic programming, but are often amenable to simulation optimization. Fu [18] reviews a host of optimization techniques, including Tabu Search, Simulated Annealing, and Genetic Algorithms, 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 implementing a Genetic Algorithm, the approach taken in this thesis, such as described in [17] and [13].

3.2.1 Ranking and Selection

Buchholz and Thümmler [7] presents a strategy for solving simulation optimization problems that utilizes a ranking and selection procedure as well as an evolutionary algorithm. In any evolutionary algorithm, such as a genetic algorithm, the algorithm must repeatedly choose the best candidate solutions from a population of candidate solutions. This type of decision is often made millions of times during a single experiment. It is vital, therefore, that these decisions be reinforced by some measure of statistical confidence. 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 based on factors such as expected means; in the literature, this probability is referred to as the *probability of correct selection*. Law [30] provides an extensive review of the problem of selecting from multiple alternatives and the different ranking and selection methods.

Rinott [38] presents a two-stage procedure to guarantee a specified lower bound on the probability of correct selection. This procedure is an improvement on the work by Dudewicz and Dalal [16] wherein a linear function of the expected means are compared; in the Rinott procedure,

ordinary expected means are compared. Both the Rinott method and the Dudewicz and Dalal method are concerned with selecting the single best alternative. 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 [29] presents 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 [9]. Ranking and selection procedures have also been extended to pareto optimization [10]. For excellent surveys on research in the field of ranking and selection, see [27], [21], and [43].

3.2.2 Genetic Algorithms

Genetic algorithms (GAs) are conceptually rooted in the evolutionary laws of natural selection, and are members of the larger class of evolutionary algorithms. Like natural selection, GAs improve a system over time by repeatedly modifying a population of candidate solutions in such a way that better candidate solutions are promoted over time. In 1975, Holland [26] proposed this natural concept as an algorithmic tool for systems modeling and analysis .

Genetic algorithms (GAs) have been widely applied across diverse disciplines. Cryptanalysts have used GAs with an initial population comprising guessed keys to attempt to decrypt intercepted cyphertext [33]. Astronomers have used GAs to fit models of the rotation curve of a distant galaxy [8]. Electrical engineers have used GAs to aid in the design of digital signal processing and analysis filters [23]. The common thread that unites these and many other applications, is the use of GAs to overcome the computational intractability of certain problems.

In the literature, certain terms have developed to emphasize the analogy to evolution and natural selection. The majority of the terms are summarized in [20]. We present the most relevant and, in our estimation, the most important of these terms. Generally, a problem’s vector of decision variables is referred to as a *chromosome*², and each individual decision variable is referred to as a *gene*. A set of *individuals* (candidate solutions) forms a *population*. Because an individual is often fully characterized by a single chromosome or string, the terms individual and chromosome are occasionally used interchangeably. A *generation* of a GA is the population present during a specific iteration of the GA; the populations of different generations often overlap. The process of forming a new generation consists of *reproduction* and *mutation*. Reproduction is the process of forming a new chromosome by some mechanism operating on two existing chromosomes. To encourage diversity and to avoid becoming trapped in local optima, a *mutation* operation is applied to the population once per generation. Mutation randomly changes a gene or a set of genes so that genetic material which may not be present in the

²There can be multiple chromosomes together constituting a *genotype*, which is the complete codification of an entity defining a solution to an optimization problem, but we omit this terminology because we find it more natural to think of a problem’s solution as a single vector of decision variables with specific values.

current population may be introduced in the following population. We refer to the quality of a particular chromosome as that chromosome's *fitness*; fitness is the measure reported by the objective function.

The concept of *elitism*, always retaining the most fit chromosome at the end of a generation, was first mentioned by De Jong [12]. Elitism has been shown to greatly improve GAs [44]. Elitism is a way of forcing some amount of fit genetic material to be passed along to the following generation without reliance upon probabilistic forces.

3.2.3 Parallel Computing

The concept of parallel computing is to gain efficiency by performing tasks simultaneously. Not every task, however, can be performed simultaneously. Tasks whose completion depend upon another task's completion, must be performed sequentially. Therefore, fundamental to parallel computing is the question of whether a particular problem can be broken down in such a way that allows for simultaneous activity. An example of a problem that can be parallelized is the multiplication of two matrices. Each component of the product matrix can be computed independently from the computation of any other component. An example of a problem that cannot be parallelized is the single simulation of a Markov chain. Since each state (except the initial state) depends upon the previous state, the simulation must be performed sequentially.

Alba and Tomassini [2] and Haupt and Haupt [24] provide comprehensive classification for parallel GAs. The first major class of parallel GAs is modelled as a *master-slave* algorithm; the master controls sorting, modifying, and program flow, whereas the slave only evaluates the fitness function. Master-slave GAs are often referred to as *panmictic* or *micro-grained* parallel GAs. These panmictic GAs can further be classified within a continuum, the extremes of which are labeled *steady-state* and *generational*. On the steady-state end, only one or two chromosomes are replaced with new chromosomes after each generation. On the generational end, the entire population of chromosomes is replaced after each generation. The second major class of parallel GAs is called *distributed* or *island* GAs. Unlike the single-population master-slave parallel GAs, distributed parallel GAs encompass multiple subpopulations which communicate to one another via migration, where a chromosome switches from one subpopulation into another. The third major class of parallel GAs is called *cellular* or *fine-grained*. For an excellent introduction to parallel GA strategies, see Goldberg [20].

3.3 Contribution to the Literature

The current work is the first attempt to optimize PSA screening policies for the detection of prostate cancer using a simulation model embedded within a metaheuristic algorithm (i.e. the GA). The methodology employed in the current work is generalizable to the problem of

optimizing treatment or screening policies for any disease where that disease may be modelled as a Markov decision process.

Simulation Model

Our simulation model is based on a non-stationary finite horizon Markov process. The health states in our model are No Cancer (no prostate cancer), Cancer (prostate cancer that is undetected), Treatment (prostate cancer detected and treatment applied), and Death (from any cause, including prostate cancer). This state definition is the same as or similar to the state definitions in other studies [46],[47],[39]. A patient's progression through these health states is modelled as a Markov process, as illustrated in Figure 4. A patient transitions between health states every 3 months and ultimately reaches *treatment* or *death*. The latter may be death from prostate cancer or death from any other cause. Once a patient reaches either treatment or death, that patient is removed from the model. In other words, these are *absorbing states*. Although the specific details of our state definitions and transition probabilities differ, it is worth noting that others have used Markov models to represent prostate cancer progression [39].

The treatment state in Figure 4 is actually an aggregate of multiple cancer states, as illustrated in Figure 4 where the rectangle represents treatment. The labels NC, C, and D represent No Cancer, Cancer, and Death, respectively. The aggregated cancer states are organ confined (OC), extraprostatic (EP), lymph node-positive (LN), and metastasis (mets). The simplification of aggregating the cancer states into the single treatment state does not result in loss of accuracy, because the reward function treats the treatment state as the expected future rewards associated with the underlying Markov reward chain [46].

The transition probabilities among NC, C, T, 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 Quality Adjusted Life Years (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 [19].




Figure 4.1: Transitions between health-states

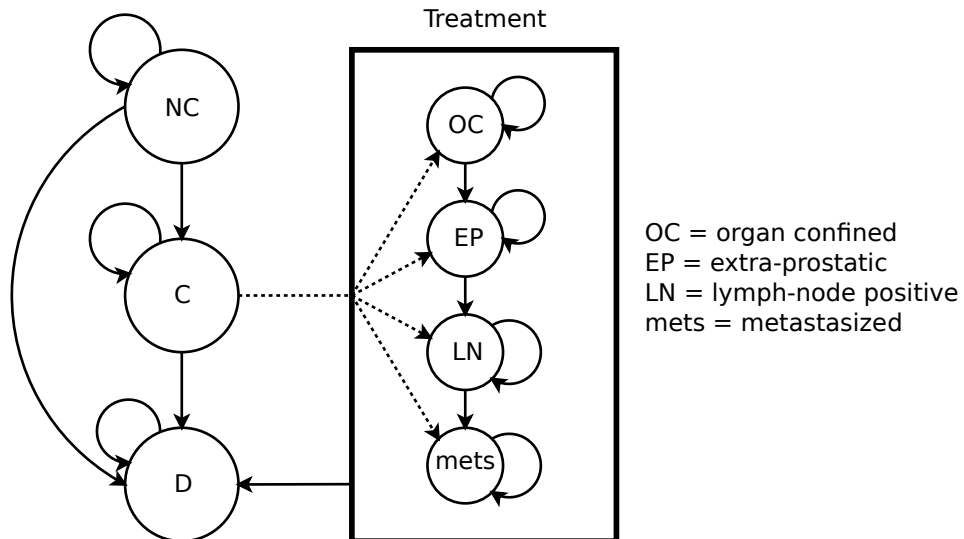


Figure 4.2: States considered following treatment

4.1 PSA Screening

Screening occurs every 3 months beginning at a defined starting age (e.g. 40 years, the age at which PSA screening is often initiated in the U.S.). At every epoch, i.e., every 3 months, a PSA level is assigned (when the screening policy calls for screening at that epoch) to the patient according to the non-stationary probability distribution of PSA levels. The distribution of PSA levels is in reality continuous. However, due to limited data, we discretize the distribution to form a more manageable approximation. Figure 4.1 illustrates the discrete PSA distribution for a patient who has not been diagnosed with prostate cancer.

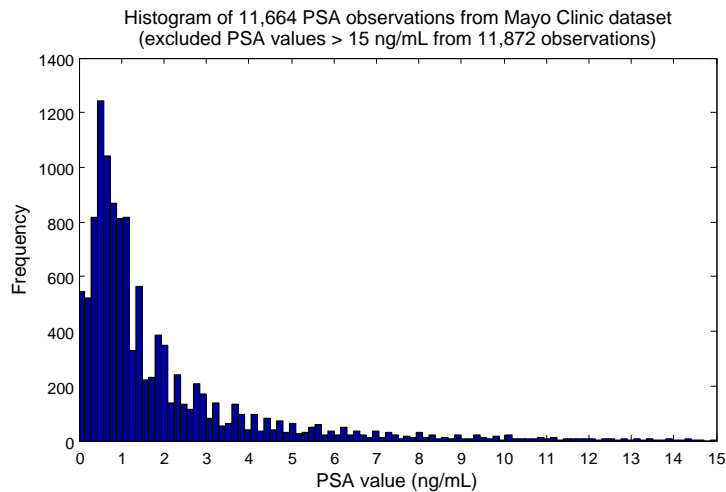


Figure 4.3: Histogram of PSA observations from empirical data

For each decision epoch, the screening policy defines whether or not to screen and, if the patient is screened, defines the criteria, i.e., PSA threshold, for a positive test result. In our model, if the patient’s PSA value is greater than the threshold value corresponding to the patient’s age (i.e. the current epoch), 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, D , as

$$D = (x_1, x_2, \dots, x_{400}),$$

where $x_i \in \{0, 0.5, \dots, 10, \infty\}$ defines a PSA threshold for biopsy referral at age i over 400 decision epochs separated by 3-month intervals (amounting to 100 years). A threshold value of ∞ is used to denote no biopsy referral, independent of the PSA test result, and therefore no

PSA test. Two additional factors may influence the decision to biopsy: screening frequency and biopsy history. We record biopsy history because doctors do not perform many repeated biopsies on a single patient; some sources suggest 3 as an upper bound on the maximum number of biopsies during a patient’s lifetime, however, our model limits patients to 1. A boolean variable is used to track whether a patient has had a biopsy, and only a limited number of biopsies are permitted during a patient’s lifetime.

4.2 Accumulation of Reward

At the beginning of each period of the Markov chain, the patient accumulates a reward which is dependent upon both the action taken and the current core health-state. In reality, a patient’s QALYs accumulate continuously; but in our discretized model, a patient accumulates QALYs at the beginning of each period. Therefore, a patient receives the *expected* reward based on the patient’s health state and action at the beginning of each period. This reward is discounted according to the patient’s age.

We abbreviate the core health-states *no cancer*, *cancer not detected*, *treatment*, and *death* as, respectively, NC, CND, T, and D. The parameters of the reward function are defined as follows: μ is the cost in QALYs of biopsy; δ is the QALYs for a patient in state T; f is the biopsy detection rate; and ϵ is the QALYs for a patient who has previously undergone treatment. We define rewards conditioned on the decision to biopsy (B) or to wait (W) according to Zhang 2009 [46], as shown below.

$$\begin{aligned} r(\text{NC}, \text{B}) &= 1 - \mu \\ r(\text{CND}, \text{B}) &= 1 - \mu - \delta(1 - f) - f\epsilon \\ r(\text{T}, \text{B}) &= 1 - \epsilon \\ r(\text{D}, \text{B}) &= 0 \end{aligned}$$

$$\begin{aligned} r(\text{NC}, \text{W}) &= 1 \\ r(\text{CND}, \text{W}) &= 1 - \delta \\ r(\text{T}, \text{W}) &= 1 - \epsilon \\ r(\text{D}, \text{W}) &= 0 \end{aligned}$$

4.3 Core Health-State Transition Probabilities

At the end of each period, the patient makes a core health-state transition. This transition occurs according to the one-step transition probability matrices below. The transition probabilities

were derived from Surveillance, Epidemiology and End Results (SEER) data and the Mayo Clinic Radical Prostatectomy Registry (MCRPR). Parameters for the transition matrices are defined in Table 4.3.

	Description
w_t	Annual prostate cancer incidence rate
d_t	Annual death rate from all causes
b_t	Annual prostate cancer death rate for state <i>Treatment</i>
e_t	Annual prostate cancer death rate for state <i>Cancer</i>
f	Biopsy detection rate

Table 4.1: Parameters for transition probability matrices

Given the decision to biopsy, we define the one-step transition probability matrix as [46]:

	NC	CND	T	D
NC	$P(\text{NC} \text{NC}, B)$	$P(\text{CND} \text{NC}, B)$	0	$P(\text{D} \text{NC}, B)$
CND	0	$P(\text{CND} \text{CND}, B)$	$P(\text{T} \text{CND}, B)$	$P(\text{D} \text{CND}, B)$
T	0	0	$P(\text{T} \text{T}, B)$	$P(\text{D} \text{T}, B)$
D	0	0	0	1

$$\begin{aligned}
P(\text{NC}|\text{NC}, B) &= (1 - d_t)(1 - w_t) \\
P(\text{CND}|\text{NC}, B) &= (1 - d_t)w_t \\
P(\text{D}|\text{NC}, B) &= d_t \\
P(\text{CND}|\text{CND}, B) &= (1 - f)(1 - d_t)(1 - e_t) \\
P(\text{T}|\text{CND}, B) &= f(1 - b_t)(1 - d_t) \\
P(\text{D}|\text{CND}, B) &= fb_t(1 - d_t) + e_t(1 - f)(1 - d_t) \\
P(\text{T}|\text{T}, B) &= (1 - d_t)(1 - b_t) \\
P(\text{D}|\text{T}, B) &= d_t + b_t(1 - d_t)
\end{aligned}$$

In the absence of biopsy (i.e. waiting), we define the one-step transition probability matrix as [46]:

	NC	CND	T	D
NC	$P(\text{NC} \text{NC}, W)$	$P(\text{CND} \text{NC}, W)$	0	$P(\text{D} \text{NC}, W)$
CND	0	$P(\text{CND} \text{CND}, W)$	0	$P(\text{D} \text{CND}, W)$
T	0	0	$P(\text{T} \text{T}, W)$	$P(\text{D} \text{T}, W)$
D	0	0	0	1

$$\begin{aligned}
P(\text{NC}|\text{NC}, W) &= (1 - d_t)(1 - w_t) \\
P(\text{CND}|\text{NC}, W) &= (1 - d_t)w_t \\
P(\text{D}|\text{NC}, W) &= d_t \\
P(\text{CND}|\text{CND}, W) &= (1 - d_t)(1 - e_t) \\
P(\text{D}|\text{CND}, W) &= d_t + e_t(1 - d_t) \\
P(\text{T}|\text{T}, W) &= (1 - d_t)(1 - b_t) \\
P(\text{D}|\text{T}, W) &= d_t + b_t(1 - d_t)
\end{aligned}$$

After making a core health-state transition, the PSA observation for the next period (if there is another) of the Markov model is observed. The simulation proceeds until the patient transitions to the absorbing state D, i.e., until the patient dies. It is important to note that the Markov chain is partially observable; the core health state is unknown to the observer, but the likelihood of cancer can be inferred from the PSA observation.

The information matrix, Q , is

$$Q = \begin{pmatrix} 0.488 & 0.318 & 0.099 & 0.058 & 0.017 & 0.020 & 0 & 0 \\ 0.253 & 0.242 & 0.135 & 0.170 & 0.082 & 0.118 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

where the rows correspond to core health-states no cancer, cancer, treatment, and death, respectively; and the columns correspond to PSA intervals $[0, 1)$, $[1, 2.5)$, $[2.5, 4)$, $[4, 7)$, $[7, 10)$, $[10, \infty)$, treatment, and death, respectively [46].

4.4 Computing Platforms

All computations for this thesis were performed on one or both of the systems below. The GNU OpenMP implementation (GOMP) was used for parallel execution. Mathworks MATLAB 7.9.0 (R2009b) was used heavily for data processing, analysis, and presentation.

System 1. Dell Linux server with 2 Quad-Core Intel Xeon E5420 2.5GHz CPUs and 16GB shared RAM. Compiler: gcc (GCC) 4.1.2 (64-bit target) by the Free Software Foundation, Inc.

System 2. IBM Blade Center Linux Cluster with 4 Quad-Core AMD Opteron 8374 HE 2.2GHz CPUs and 64 GB shared RAM per blade. Compiler: pgCC 7.2-2 (32-bit target) by The Portland Group, Inc.

4.5 Numerical Results

4.5.1 Data and Model Parameters

Table 4.5.1 list the values and sources for the reward parameters used in evaluating the proposed screening policies, where λ is the discount factor for the reward function, ϵ is the annual utility decrement of living in state T, μ is the one-time utility decrement associated with biopsy, δ is the annual utility decrement of living in state C, and f is the biopsy detection rate for patients with prostate cancer. Unless explicitly stated otherwise, all results in this thesis pertain to the set of parameter values in Table 4.5.1.

Parameter	Value	Source
ϵ	0.05	Bremner et al. 2007
μ	0.01	Chhatwal et al. 2008
δ	0	Bremner et al. 2007
λ	1	Gold et al. 1996
f	0.8	Haas et al. 2007

Table 4.2: Reward parameters used in evaluating proposed screening policies

Table 4.5.1 list the values and sources for the state transition probabilities used in evaluating the proposed screening policies, where w_t is the annual prostate cancer incidence rate, d_t is the annual death rate from other causes, b_t is the annual prostate cancer death rate for patients in state T, e_t is the annual prostate cancer death rate for patients in state C, and f is the biopsy detection rate for patients with prostate cancer. For Table 4.5.1, note that MC-RPR stands for the Mayo Clinic Radical Prostatectomy Registry.

4.5.2 Discounting QALYs

The parameter $\lambda = 1$ implies there is no discount over time associated with QALYs, whereas, the QALYs are discounted over time in the literature [19]. However, discounting life years

Parameter	Value	Source
w_t	Age Specific	National Cancer Institute 2008
d_t	Age Specific	National Cancer Institute 2008, Heron 2007
b_t	Age Specific	MC-RPR
e_t	0.033	Albertsen et al. 2005
f	0.8	Haas et al. 2007

Table 4.3: Transition probabilities used in evaluating proposed screening policies

is controversial, and has a significant effect on optimal policies when rewards are discounted over large time intervals (e.g. decades of a patient’s life). Therefore, we let $\lambda = 1$ in order to better distinguish the policies, as illustrated in Figure 4.5.2. Although Figure 4.5.2 appears to graph 6 different functions, it actually graphs 18 different functions; each line is formed by the combination of 3 different lines, where $\mu = 0.01, 0.05, \text{ and } 0.1$. At this level of magnification, functions differing only by μ values in the range $[0.01, 0.1]$ overlap and are indistinguishable.

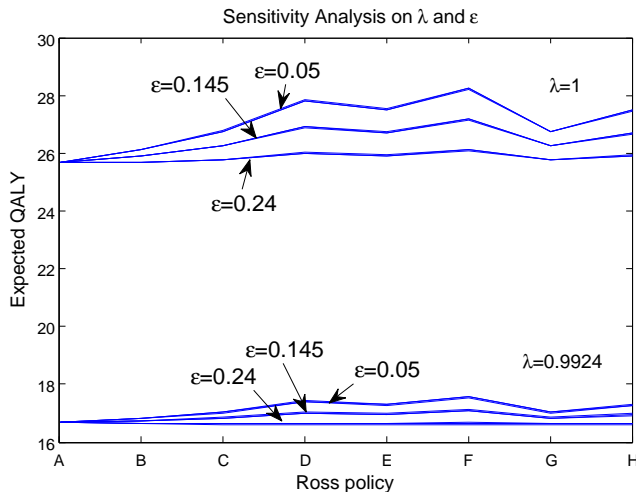


Figure 4.4: Sensitivity analysis on parameters of reward function for Ross policies

4.5.3 Evaluation of Proposed Screening Policies

A Monte-Carlo simulation by Ross (2007) compared several prostate-cancer screening policies using as outcome measures the ratio of the number of prostate cancer deaths prevented to the number of PSA tests and biopsies performed. We now compare the same screening policies using

the outcome measure quality-adjusted life years (QALYs). These policies (hereafter referred to as the *Ross policies*) are defined as follows: (A) no screening; (B) threshold 4.0 ng/mL tested every 5 years from age 40 to 75; (C) threshold 4.0 ng/mL tested every 2 years from age 50 to 75; (D) annual testing with threshold 3.5 ng/mL from age 50 to 59, 4.5 ng/mL from age 60 to 69, and 6.5 ng/mL from age 70 to 75; (E) threshold 4.0 ng/mL tested annually from age 50 to 75; (F) threshold 2.5 ng/mL tested annually from age 50 to 75; (G) threshold 4.0 ng/mL tested at age 40, age 45, and every 2 years from age 50 to 75; and (H) threshold 4.0 ng/mL tested at age 40, age 45, and annually from age 50 to 75.

In Figures 4.5.3 and 4.5.3, we examine the effect of the PSA screening policies from Ross (2007) in terms of QALYs. Figure 4.5.3 illustrates results for the entire US male population, whereas Figure 4.5.3 only treats the subpopulation of males who ultimately develop cancer between the ages 50 and 75. It is clear from visual inspection that the range of expected QALYs in Figure 4.5.3 is small. This finding is central to the debate on the benefits of PSA screening. However, the range of expected QALYs in Figure 4.5.3 is much larger. Thus, the policies are significantly different from the perspective of their impact on the most-affected subpopulation. Figure 4.5.3 shows the importance of relying on the most-affected subpopulation — as opposed to the entire US male population — to evaluate a policy on the basis of QALYs.

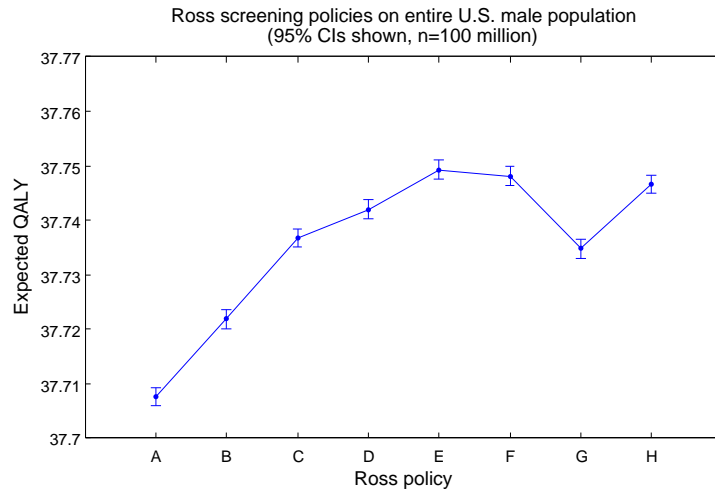


Figure 4.5: Ross results comparison for entire U.S. male population

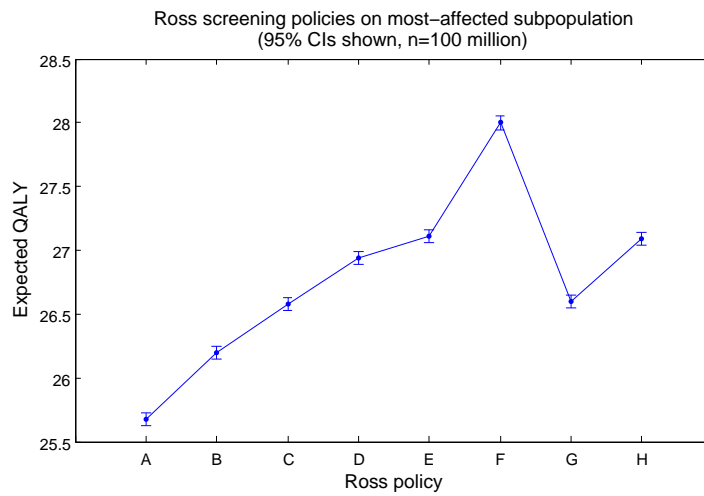


Figure 4.6: Ross results comparison for subpopulation of U.S. males who ultimately develop prostate cancer

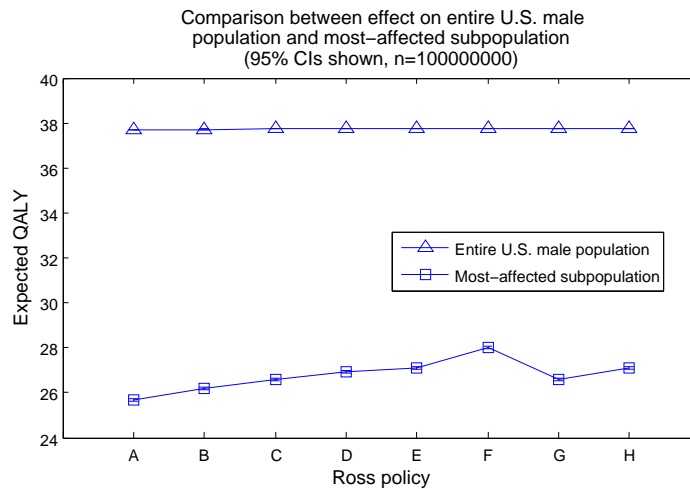


Figure 4.7: Comparison between the Ross policies' effect on the entire U.S. male population and the subpopulation of U.S. males who ultimately develop prostate cancer

4.6 Conclusions

We provide to our knowledge the first results comparing PSA threshold policies based on QALYs. The results from simulating the Ross policies reveal that the best screening policy from the perspective of the general male population is to screen annually from the age of 50 to 75 years with a threshold of 4.0 ng/mL. This policy is the traditional guidelines for PSA-threshold screening currently in practice in the U.S.. Therefore, according to the metric QALYs, and with respect to the entire population of U.S. males, the current PSA screening policy in practice in the U.S. is better than the other Ross policies, such as a common European practice of annual screening from the age of 50 to 75 years with a threshold of 2.5 ng/mL. From a practical standpoint, however, the difference between Ross policies E and F (i.e. between the 4.0 or 2.5 ng/mL thresholds) may be deemed negligible by some.

There are, of course, limitations in our model. Several model parameters are estimates (e.g. μ , ϵ) that differ based on individual patient preferences. Our results are based on a single U.S. population which may have better access to healthcare than the general population, and which is primarily caucasian. There is bias in our probability distribution for PSA levels. Specifically, we assume people without a positive biopsy in our original dataset do not have cancer, although they may in fact have cancer which has not yet been detected.

Optimization of PSA-Threshold Screening Policies

In this chapter we discuss our simulation-optimization model and solution method for prostate cancer screening. The optimization problem is to find the set of PSA thresholds as a function of age that maximizes the expected QALYs over a patient's lifetime.

Our genetic algorithm (GA) explores the space of possible screening policies and evaluates each screening policy by simulation at a specified degree of precision. The greater the portion of the policy space explored by the GA, *ceteris paribus*, the better the GA. The greater the emphasis on highly fit individuals, *ceteris paribus*, the better the GA. By increasing exploration, we have a higher probability of reaching the global optimum. By increasing exploitation, we have a higher probability of increasing fitness from generation to generation. Unfortunately, these two attributes—the ability to search a large portion of the policy space (exploration) and the emphasis on promoting higher fitness individuals (exploitation)—are competing; to increase one, is to decrease the other. Therefore, a fundamental consideration in implementing our GA is to find an optimal balance between exploration and exploitation.

Researchers have proposed various strategies for achieving the optimal balance [45],[3], and [35]. Finding the optimal balance between the creation of diversity and the reduction of diversity in favor of exploiting higher fitness individuals is critical to the successful performance of any GA [5]. For more on the tradeoff between exploration and exploitation, see Goldberg [20] or Spears [42].

To address the exploitation/exploration dilemma, we use a combination of methods including the simulation model of Chapter 4, a GA, and a well known ranking and selection method. We evaluate different variants of the GA and examine the corresponding fitness levels at which the

variants converge as well as how rapidly the variants reach convergence. We experiment with the ranking and selection to devote just enough computational effort to achieve the desired level of exploitation. Chromosomes are evaluated simultaneously on multiple processing cores to make efficient use of computing time.

The remainder of this chapter is structured as follows. After formally defining the optimization problem, we discuss how the ranking and selection procedure gives us a statistical guarantee that the GA will tend toward fitter policies, and how our ranking and selection procedure in concert with parallel processing lightens the computational burden. We then describe in detail the GA and present pseudo-code that may be adapted to other disease treatment/screening optimization problems. Next, we present sensitivity analysis on characteristics of the GA itself (e.g. mutation rate, crossover method). Finally, we present sensitivity analysis on prostate cancer screening policies with respect to parameters of the reward function and transition probabilities.

5.1 Optimization Model

For the PSA screening optimization problem, we seek a PSA screening policy represented by the vector of decision variables, D^* , such that

$$D^* = \arg \max_{D \in \mathcal{D}} J(D), \quad (5.1)$$

where

$$J(D) = E_D \left[\sum_{t=t_0}^T \lambda^t r(s, a) \right]$$

is the objective function measuring expected QALYs as presented in Chapter 4; $\mathcal{D} = \langle d_1 \times d_2 \times \dots \times d_N \rangle$ where $d_i \in \{0, 0.5, \dots, 10, \infty\}$ defines the feasible set of PSA-threshold screening policies; λ is the discount factor over time t ; 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 $r(s, a)$ is the reward function as a function of the health state and action.

The optimization problem is difficult because of the combinatorially large number of possible PSA-threshold screening policies and the lack of exploitable mathematical properties (e.g. convexity). Given the limitation of discrete PSA thresholds as defined in the previous paragraph, and assuming $t_0 = 160$ (40 years), there are $22^{240} = 1.5 \times 10^{322}$ possible policies. Each policy requires a non-trivial amount of time to evaluate, and in reality the range of PSA thresholds is continuous. Solving this problem by exhaustive enumeration is infeasible. Therefore, the most appropriate way to address this problem is to use a metaheuristic search algorithm.

5.2 Genetic Algorithm

Genetic algorithms (GAs) search the space of feasible screening policies and evaluate the fitness of the screening policies by simulation. In our simulation-optimization approach, the simulator is embedded in the GA and serves to provide statistical estimates of the expected value of any candidate optimal policy. A high-level model of the algorithm is shown in Figure 5.2.

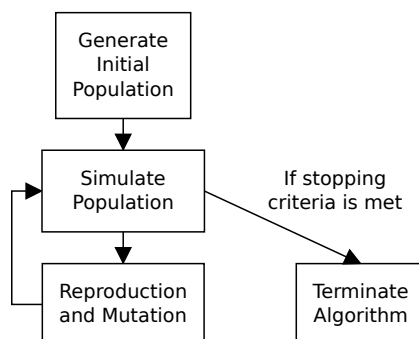


Figure 5.1: High-level model of GA

We use ranking and selection to determine the appropriate sample size for simulating individuals in the population. This provides a statistical guarantee on the probability of correct selection, $P(CS)$. In particular, we use the *Rinott procedure*. Rinott (1978) presents a now commonly-used two-stage procedure to select the single best of several alternatives [38]. The Rinott procedure is suited to our algorithm because it selects the single best individual (chromosome) from the population that is to be passed along unmodified to the next generation in keeping with the concept of elitism; a similar strategy is discussed in [7].

The Rinott procedure is based on the *least favorable condition*, where the best individual differs from the remaining individuals by the *indifference amount*, δ^* , and the remaining individuals each differ one-to-another by less than the indifference amount. The indifference amount, which is highly problem specific, is the amount greater than or equal to which two individuals must differ in order to be regarded as significantly different. In mathematical terms, the least favorable condition is an arrangement of each individual, D_i , such that $J(D_i) = J(D^*) - \delta^*$ for all $D_i \neq D^*$. Because the procedure is based on this least favorable condition (which rarely occurs), it achieves a lower bound on the $P(CS)$. The parameters and constants

used are defined in the following integral equation:

$$\int_0^\infty \left[\int_0^\infty \Phi \left(\frac{h}{\{(N_0 - 1)[(1/x) + (1/y)]\}^{1/2}} \right) f(x) dx \right]^{k-1} f(y) dy = P^*, \quad (5.2)$$

where P^* is a lower bound on the probability of correct selection, N_0 is the stage-one sample size for variance estimation, k is the total number of independent normal individuals in the population, x and y are independent Chi-squared random variables with $N_0 - 1$ degrees of freedom, Φ is the CDF for the standard normal distribution, h is a constant, and $f(x)$ and $f(y)$ are the density functions of x and y , respectively. The steps of the Rinott procedure are:

1. Simulate each individual D_i using sample size N_0
2. Simulate each individual D_i using sample size $N_i - N_0$, where $N_i = \max \left\{ N_0, \left(\frac{h}{\delta^*} \right)^2 S_{D_i}^2 \right\}$, δ^* is the indifference-zone amount, and $S_{D_i}^2$ is the variance of the set of N_0 simulations of D_i
3. Compute the expected QALYs for each individual D_i , where $\bar{X}_i = \frac{1}{N_i} \sum_{j=1}^{N_i} J(D_{ij})$ is the expected QALYs of D_i , and the random variable X is the simulation performance measure in QALYs
4. Select as best the individual D_i with largest \bar{X}_i

The basic idea of the Rinott procedure is to provide a guarantee that the single best individual from a population will be selected at least a chosen probability of the time. The procedure allows us to specify a lower-bound on the probability at which we select the best policy from a set of policies. The fundamental information that the procedure provides to us is the sample size that is sufficient and no larger than necessary to guarantee the chosen probability of correct selection.

At each generation of the GA, we want the single best policy from the population to persist unchanged unto the next generation. The Rinott procedure allows us to do exactly that without wasting valuable computing time. Therefore, this procedure is an appropriate addition to our optimization algorithm. In addition to providing a lower-bound on 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 the chromosome selection process. This is due, mathematically, to the fact that the solution to the integral in (5.2) for our problem with $k = 2$ is greater than the solution for the same problem with $k > 2$.

In order to select individual chromosomes to mate and produce offspring for the subsequent generation, we employ and compare two standard methods. The first is the *Roulette Wheel*

selection algorithm, which the authors in [48] claim to be the most frequently used selection algorithm. The Roulette Wheel algorithm is a type of fitness-proportional selection method. The probability that a particular chromosome is chosen to be a parent is directly related to the chromosome's fitness relative to the combined fitness of the entire population.

Each chromosome, k , is ordered by and assigned to a section of the wheel proportional to the chromosome's fitness (see Figure 5.2). A random number uniformly distributed on the interval $[0,1]$ is generated. The value of the random number will fall into a section of the wheel. The chromosome into whose section of the wheel the random number falls, is the chromosome selected with replacement for reproduction. This selection process repeats until enough chromosomes are selected to produce the required offspring. The offspring will tend to have relatively high fitness values, and, therefore, the offspring is biased towards better prostate-cancer screening policies.




Figure 5.2: Example roulette wheel illustration

The second approach we employ is *Tournament* selection, which some [48] claim to be often superior to the Roulette Wheel algorithm. Tournament selection is not a fitness-proportional method like the Roulette Wheel approach. In Tournament selection, two chromosomes are uniformly randomly drawn from the population, and the chromosome with greater fitness is selected to become a parent for reproduction. This process repeats until enough parents are selected to create the next generation of chromosomes. At each round of tournament selection, all chromosomes are available; in other words, the chromosomes are randomly drawn *with replacement* from the population. Our implementation of tournament selection is technically referred to as *deterministic* tournament selection, because at each pairwise comparison, the

chromosome with greater fitness is selected with probability 1.

The selection method describes the process by which a chromosome is selected to have its genetic material promoted to subsequent generations. Specifically, the selection method is used to determine parents for the reproduction (crossover) stage. We examine the performance of the GA under the two selection methods, the Roulette Wheel method, and the Tournament method.

The process of mating two parent chromosomes to produce a new child chromosome is accomplished by crossover. The idea behind crossover is to form a new chromosome that shares features of two existing chromosomes. In our problem, the two parents are defined by the following vectors of decision variables:

$$\begin{aligned} D_{P1} &= (l_1, l_2, \dots, l_N) \\ D_{P2} &= (r_1, r_2, \dots, r_N). \end{aligned}$$

We compare three different crossover methods. The first is called *Single-Point* crossover [44], and is illustrated by Figure 5.2. This algorithm randomly selects a location q , where $q \in \mathbb{Z} : 0 < q < N$. The vector of decision variables of the resulting child, D_C , acquires the first q decision variables from the first parent, D_{P1} , and the remaining $N - q$ decision variables from the second parent, D_{P2} , so that the child, after crossover, is defined as

$$D_C = (l_0, \dots, l_{q-1}, r_q, \dots, r_N).$$

The second crossover method is called *Two-Point* crossover [44], and is illustrated by Figure 5.2. This algorithm randomly selects two locations p and q . First, p is randomly selected such that $p \in \mathbb{Z} : 0 < p < N - 1$. Then, q is randomly selected such that $q \in \mathbb{Z} : p < q < N$. The vector decision variable of the resulting child, D_C , acquires the first $0 \dots p$ decision variables from the first parent, D_{P1} , the middle $p + 1 \dots q$ decision variables from the second parent, D_{P2} , and the last $q + 1 \dots N$ decision variables from the first parent, D_{P1} . The resulting child is, therefore, defined as

$$D_C = (l_0, \dots, l_p, r_{p+1}, \dots, r_q, l_{q+1}, \dots, l_N).$$

The third crossover method is called *Arithmetic* crossover [25]. This algorithm simply takes the means (rounded to nearest integers) of the decision variables of both parents, D_{P1} and D_{P2} , as the decision variables of the child, D_C . The resulting child is, therefore, defined as

$$D_C = (\mathcal{F}(l_0, r_0), \dots, \mathcal{F}(l_N, r_N)),$$

where the function $\mathcal{F}(a, b)$ returns the rounded average of a and b .

Although the mechanics of these crossover methods differ, they all serve to mimic the




Figure 5.3: Single-point crossover illustration




Figure 5.4: Two-point crossover illustration

inheritance of traits that occurs when two species reproduce in nature. The product of this reproduction (i.e. the child) is a combination of the genetic material of the parents. In this way, the hope is to pass good genetic material from two individuals onto a new individual which will replace a less fit member of the population. After time, good genes are promoted over bad genes, and the fitter genetic material survives with greater likelihood.

The crossover method defines the mechanism by which the genetic material of two chromosomes are combined in order to create a new chromosome; in other words, crossover defines the mechanics of reproduction. We examine the performance of the GA under each of the 3 crossover methods implemented.

The mutation rate is the probability that a gene (i.e. decision variable) will be mutated (i.e. randomly changed). The mutation rate is equal for every gene. The goal in selecting a mutation rate is to choose the rate which satisfactorily promotes diversity without disturbing substrings that positively contribute to the fitness of the chromosomes. From Figure 5.3.1, we conclude that the mutation rate 0.001 is appropriate for our GA; accordingly, the mutation rate 0.001 is often recommended in the literature [44].

The genetic algorithm provides a heuristic to estimate D^* . Since it is not feasible to solve (5.1)

optimally, the GA provides a lower bound on the maximum expected QALYs that can be achieved through PSA-threshold based prostate cancer screening. Pseudo-code for the genetic algorithm is shown in Algorithm 1; the genetic algorithm's sub-routines are shown in Algorithms 2, 3, and 4.

Algorithm 1: GENETICALGORITHM()

```

1 begin
2   Generate initial population of policies;
3   while stopping criteria not met do
4     foreach policy  $D_i$  do
5       Use Rinott method to determine sample size,  $N_i$ , for desired  $P^*$ ;
6       Simulate policy using determined sample size;
7       Record  $J(D_i)$ ;
8     BESTPOLICY  $\leftarrow$  policy with greatest expected QALYs;
9      $J^{\max} \leftarrow J(\text{BESTPOLICY})$ ;
10    foreach policy  $D_i$  do
11      if  $J(D_i) \neq J^{\max}$  then
12        PARENT1  $\leftarrow$  TOURNAMENTSELECT();
13        PARENT2  $\leftarrow$  TOURNAMENTSELECT();
14        NEWPOLICY  $\leftarrow$  CROSSOVER(PARENT1, PARENT2);
15        MUTATE(NEWPOLICY);
16         $D_i \leftarrow$  NEWPOLICY;
17    return BESTPOLICY;
18 end

```

Algorithm 2: TOURNAMENTSELECT()

Output: Selected policy

```

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

```

Algorithm 3: CROSSOVER(Parent1, Parent2)

Input: Parent policies which reproduce to form a new policy

Output: New policy formed by reproduction

```
1 begin
2    $cxpoint1 \leftarrow$  random integer  $\in [1, 400 - 1]$ ;
3    $cxpoint2 \leftarrow$  random integer  $\in [cxpoint1 + 1, 400]$ ;
4   for  $i = 1$  to  $cxpoint1$  do
5      $\lfloor$  NewPolicy[i]  $\leftarrow$  Parent1[i];
6   for  $i = cxpoint1 + 1$  to  $cxpoint2$  do
7      $\lfloor$  NewPolicy[i]  $\leftarrow$  Parent2[i];
8   for  $i = cxpoint2 + 1$  to 400 do
9      $\lfloor$  NewPolicy[i]  $\leftarrow$  Parent1[i];
10  return NewPolicy;
11 end
```

Algorithm 4: MUTATE(Policy)

Input: Policy to be mutated

Output: Mutated policy

```
1 begin
2   for  $i = 1$  to 400 do
3      $x \leftarrow$  uniform random variate  $\in [0, 1]$ ;
4     if  $x \leq 0.001$  then
5        $\lfloor$  Policy[i]  $\leftarrow$  randomly selected PSA threshold  $\in \{0, 0.5, \dots, 10, \infty\}$ ;
6   return Policy;
7 end
```

5.3 Results

The following subsections contain sensitivity analysis on running time of the simulator (objective function estimator), sensitivity analysis on parameters of the reward function, sensitivity analysis on parameters of the GA, and numerical analysis of the optimal policy produced by the GA.

5.3.1 Experimental Evaluation of Genetic Algorithm

Table 5.3.1 provides results from single simulations run on a single threshold and screening frequency policy, where random variates in the stochastic model were drawn from the same random number stream. These results show the effect of sample size and point out the tradeoff between precision of our expected QALY calculation at the 95% confidence level and the compute time required to make such calculation. The relative precision is the confidence interval half-length divided by the midpoint of the confidence interval; this measure expresses the margin for error as a fraction of the estimated quantity (i.e. expected QALYs).

#	Sample size	Half width (QALY)	Relative precision	Run time (sec)
1	1,000	22.615	94.58%	0.02
2	10,000	4.972	16.28%	0.24
3	100,000	1.530	5.65%	2.38
4	200,000	1.091	4.06%	4.77
5	300,000	0.873	3.32%	7.15
6	400,000	0.754	2.87%	9.53
7	500,000	0.671	2.56%	11.90
8	600,000	0.615	2.35%	14.29
9	700,000	0.566	2.16%	16.67
10	800,000	0.525	2.01%	19.06
11	900,000	0.495	1.89%	21.42
12	1,000,000	0.471	1.80%	23.81
13	10,000,000	0.151	0.57%	238.13
14	100,000,000	0.048	0.18%	2392.50

Table 5.1: Precision and running time analysis

Figure 5.3.1 shows the expected QALYs as a function of sample size. By this plot, it is clear that sample sizes less than 300,000 will produce a mean which has not yet settled to the degree necessary to make distinctions between expected QALYs differing by 1. The goal in selecting a sample size for the *exploration* stages in any metaheuristic is to find a suitable tradeoff between time and accuracy. When the metaheuristic initially evaluates populations in order to rank

them and produce *better* offspring, all that is needed is a reliable measure to determine the optimality of the populations relative to each other. At this stage, we do not need to estimate the actual optimal value of the populations; a simple ordering is sufficient. Since smaller sample sizes require less compute-time, selecting a sample size which is no more precise than necessary allows the metaheuristic to search a larger solution space in a given period of time.

It is notable that a very large number of samples is required to achieve high precision. This is consistent with findings in the medical literature that demonstrate difficulty in differentiating screening policies through randomized controlled trials. This also motivates the need to carefully consider computational effort and the tradeoff between exploration and exploitation, as in our proposed simulation-optimization method.

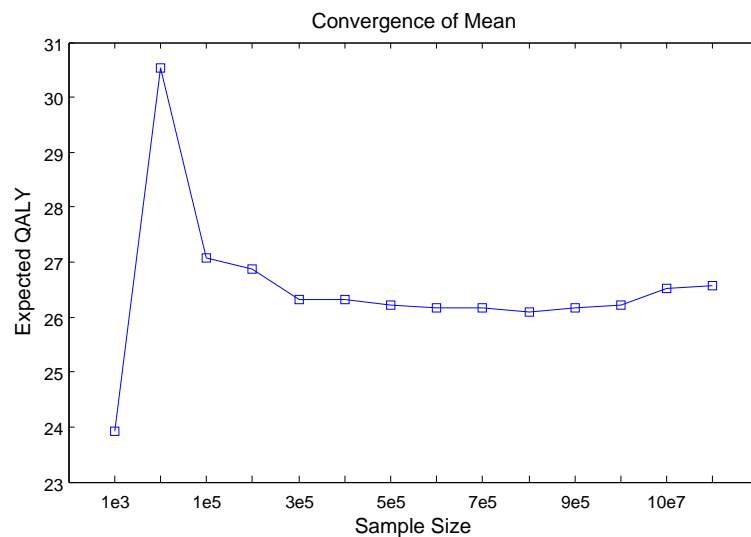


Figure 5.5: Convergence of mean as sample size increases

The performance of the single-point, two-point, and arithmetic crossover methods are shown below, respectively, in Figure 5.3.1. Although the single-point method and the two-point method achieve roughly the same level of fitness at the end of 500 generations, the two-point crossover method shows slightly more rapid convergence. For this reason, we prefer the two-point method over the single-point method. By inspection it is clear that the arithmetic method performs poorly in comparison to the single-point and the two-point methods. On the basis of this analysis, we conclude that the two-point crossover method is superior for our particular problem.

From Figure 5.3.1 we see that tournament selection performs significantly better than roulette selection. If there is little difference in fitness between each chromosome in the population,




Figure 5.6: Comparison of crossover methods

then each chromosome will have nearly equal relative fitness. Because roulette selection uses the relative fitness of each chromosome to determine the probability of the chromosome being selected, each chromosome is nearly equally likely to be selected. Therefore, when there is little difference in fitness between each chromosome, roulette selection is nothing more than uniform random selection. Tournament selection, on the other hand, provides some assurance that more fit chromosomes will be selected over less fit chromosomes. Unlike roulette selection, tournament selection does not regard the magnitude of differences in fitness. In the PSA screening problem, the differences in fitness between chromosomes are slight. Therefore, we should intuitively expect Tournament selection to perform better.




Figure 5.7: Comparison of selection methods

There is no obvious victor in the mutation rate comparison shown in Figure 5.3.1. Baluja and Caruana [6] compared the mutation rate of the standard GA by varying the rate from 0.001 to 0.010, and found that 0.001 performed slightly better than 0.010. Similarly, we take 0.001 as the lower bound on the mutation rate for the PSA screening problem. From Figure 5.3.1, the

slope of the fitness function decreases sooner and more rapidly as the mutation rate increases from 0.001. Therefore, we take 0.001 as the mutation rate for the PSA screening problem.




Figure 5.8: Comparison of mutation rates

5.3.2 Experimental Evolution of Genetic Algorithm

Figures 5.3.2 and 5.3.2 display the evolutionary progress (i.e. the results over time) of the GA using two-point crossover, tournament selection, and mutation rate 0.001; only the most fit chromosome at each generation is used in the graphs. It is clear that the fitness progresses over time. At the 95% confidence level, the fitness of the initial population is statistically different from the fitness of the populations beginning around the 200th generation. Therefore, we have statistical confidence that the most-fit policy improves over time.

From Figure 5.3.2 we see that the GA appears to converge to an optimum around the 1,500th generation. Because GAs are heuristic search algorithms, there can be no guarantee derived from the GA itself that the global optimum has been reached; it is always possible to converge to a local optimum. However, we believe a near optimal solution has been reached (we provide justification for this in Section 5.3.5).

5.3.3 Optimal PSA Screening Policies

Figure 5.3.3 displays the raw data for the optimal policy generated by the GA. The random behavior seen to the right of age 65-70 is inconsequential because in the model used to compute the optimal policy, only 1 biopsy is permitted during a patient’s lifetime. Because the threshold is 0 between 65-70 years, every patient who has not yet received a biopsy receives a biopsy around

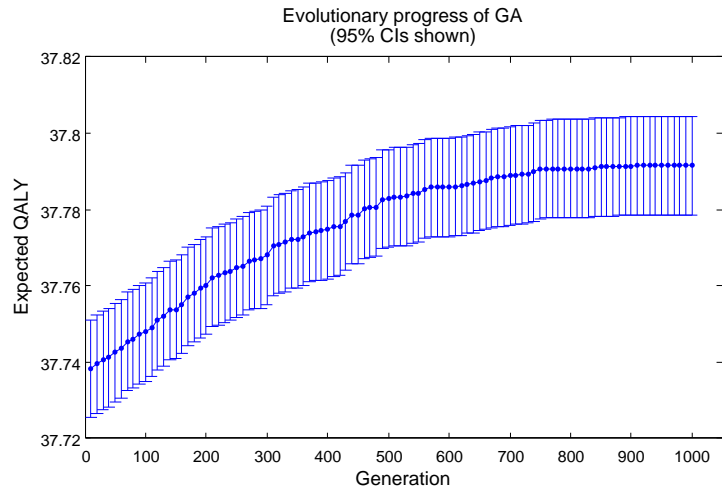


Figure 5.9: Evolution of GA showing period of sharpest incline in fitness improvement from generation to generation

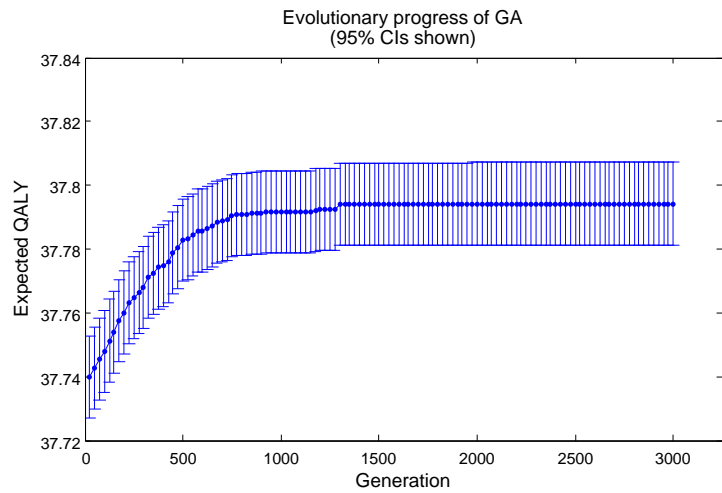


Figure 5.10: Evolution of GA showing convergence to optimum

the age 65-70. Once a patient is biopsied, the patient is removed from the model. Therefore, all data after the first 0.0 ng/mL threshold is meaningless, and the 0.0 ng/mL threshold serves the purpose of excluding screening beyond the age 65-70.



Figure 5.11: Optimal PSA screening policy discovered by GA in raw form

5.3.4 Clinically Relevant Screening Policies

Because of the nature of a GA, the optimal policy discovered by the GA will inevitably exhibit some imperfect or non-smooth patterns. In the context of a clinically implementable screening policy, these imperfectly patterned policies should ideally be *simplified* to a policy that is reasonable to implement. We consider two simplification methods. First, for the policy shown in Figure 5.3.4, starting from the initial age of screening we accept only those thresholds lying on a non-decreasing curve. That is, for each decision epoch, t_n , if the PSA threshold at t_{n-1} is less than or equal to the PSA threshold at t_n , we accept the threshold at t_n ; otherwise, we discard the threshold at t_n . (Note that we also arbitrarily discarded as an anomaly the threshold at $t_0 =$ age 40.) Second, for the policy shown in Figure 5.3.6, we arbitrarily accept certain thresholds such that the simplified policy appears more reasonable to implement. This second method does not lend itself to algorithmic description because it is inherently ad hoc and arbitrary.

There is no mathematical *motivation* for using such an ad hoc method, however, there is mathematical *justification*. Ultimately, the goal is to form a simplified policy that is nearly

as good as the raw policy, but more reasonable to implement. Because, in both Figure 5.3.4 and Figure 5.3.6, the 95% confidence intervals for the raw and simplified optimal policies overlap, or nearly overlap, we cannot say that the simplified policies are better or worse than the respective raw policies. By inspection, it is clear that the simplified policies—though not trivial to implement—are indeed more reasonable to implement than the respective raw policies. Having achieved these goals, the simplified policies and the methods of simplification are valid.



Figure 5.12: Optimal PSA screening policy discovered by GA in simplified form

Judging from Figure 5.3.4, it is clear that the best policy performs better than the Ross policies from the perspective of the entire US male population. Judging from Figure 5.3.4, however, the optimal policy is *not* optimal with respect to the subpopulation of males who ultimately develop cancer. The improvement in Figure 5.3.4 is greater than the loss in Figure 5.3.4; this type of comparison is central to the optimization of screening policies for low probability diseases like prostate cancer. Very rarely will a single screening policy for a disease be optimal with respect to all subpopulations—especially with respect to a general population and a most-affected subpopulation. And the optimal policy in Figure 5.3.4 is on par with most of the Ross policies. For these reasons, we believe the best policy generated by the GA to be close to the overall optimal PSA screening policy.

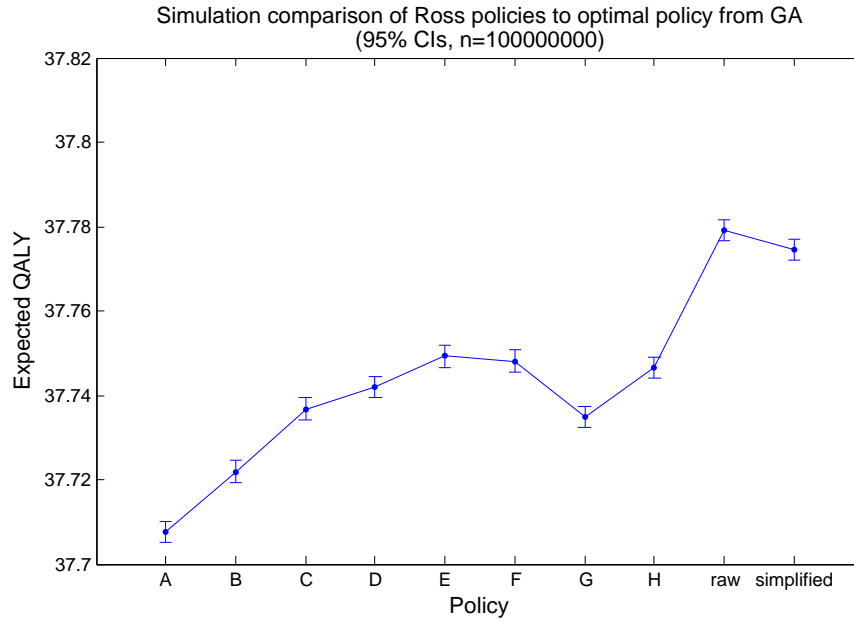


Figure 5.13: Comparing the Ross policies to the optimal GA policy (both the raw version and the simplified version) in terms of the effect on the entire U.S. male population

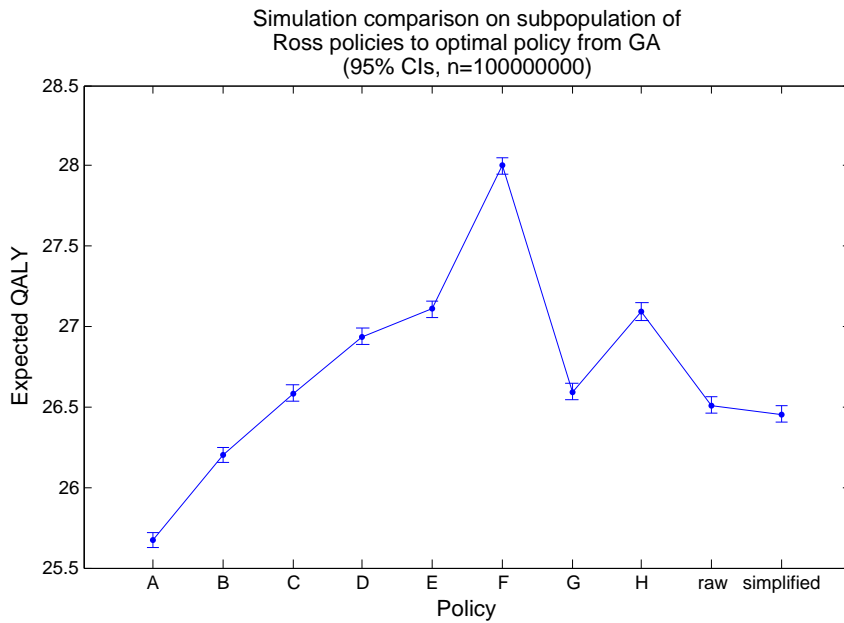


Figure 5.14: Comparing the Ross policies to the optimal GA policy in terms of the effect on the subpopulation of U.S. males who ultimately develop prostate cancer

5.3.5 Upper Bounds on the Optimal Policy

Zhang (2010) formulated a partially-observable Markov decision process (POMDP) to determine the theoretical optimal PSA-threshold screening policy using the same original data set that we used. The POMDP uses Bayesian updating to use all prior PSA tests to estimate the patient’s probability of having prostate cancer. This probability is then used to decide whether to refer the patient for biopsy and subsequent treatment. This is markedly different from our GA, which specifies according to age the PSA threshold at which to biopsy based on a single independent PSA observation.

In Table 5.2, we compare the benefits of the optimal policy found using the POMDP in Zhang (2010) with the benefits of the near optimal policy found using our genetic algorithm (shown in Figure 5.3.4). The POMDP model and the best GA policy found for comparison to the POMDP model, use parameter values $\epsilon = 0.145$ and $\mu = 0.05$; the other parameter values are as previously defined in Table 4.5.1. Because the POMDP provides theoretical upper bounds on the optimal policy, it is expected that the best GA policy is less effective. However, it is worth noting that PSA-based thresholds are the norm in clinical practice. Whereas the POMDP solution provides insights into the theoretical maximum benefit in QALYs of PSA screening, the GA provides insights into how close we may come in practice to the theoretical maximum benefit.

	Per person expected improvement (in QALYs) over no screening	Per person expected improvement (in QALYs) over traditional guidelines
GA	0.008	0.022
POMDP	0.033	0.060

Table 5.2: Comparison between the best GA policy and the optimal POMDP policy [47]

5.3.6 Multiple Biopsy

The assumption that every patient may undergo at most *one* biopsy during his lifetime is reasonable. In our dataset 81.4% of patients biopsied underwent only 1 biopsy. However, it is worthwhile to investigate the best GA policy under the assumption that every patient may undergo at most *three* biopsies. The limit of 3 biopsies is supported by anecdotal clinical evidence.

We now present and analyze the best PSA screening policy discovered by the GA under the

assumption that every patient may undergo up to three biopsies. Figures 5.3.6 and 5.3.6 show the best policy found by the GA under the three biopsy assumption in raw and simplified form, respectively.




Figure 5.15: Optimal PSA screening policy under maximum of 3 biopsies per patient discovered by GA in raw form

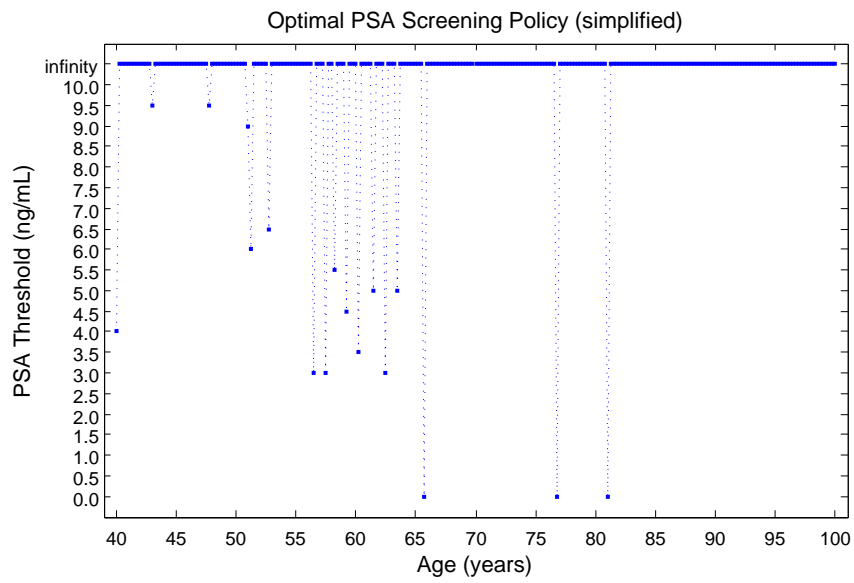


Figure 5.16: Optimal PSA screening policy under maximum of 3 biopsies per patient discovered by GA in simplified form

Judging from Figure 5.3.6, it is clear that the best policy performs better than the Ross policies from the perspective of the entire US male population. Judging from Figure 5.3.6, however, the best policy is *not* optimal with respect to the subpopulation of males who ultimately develop cancer. The improvement in Figure 5.3.6 is greater than the loss in Figure 5.3.6. And the optimal policy in Figure 5.3.6 is on par with most of the Ross policies. For these reasons, we believe the best policy generated by the GA to be the overall optimal PSA screening policy.

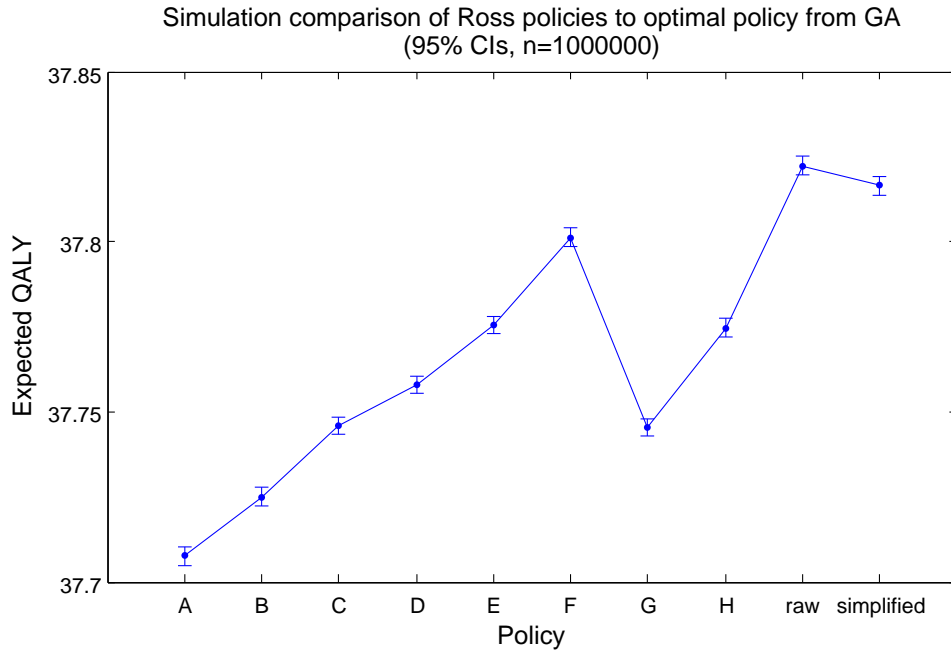


Figure 5.17: Comparing the Ross policies, under maximum of 3 biopsies per patient, to the optimal GA policy (both the raw version and the simplified version) in terms of the effect on the entire U.S. male population




Figure 5.18: Comparing the Ross policies, under maximum of 3 biopsies per patient, to the optimal GA policy in terms of the effect on the subpopulation of U.S. males who ultimately develop prostate cancer

5.3.7 Sensitivity Analysis

We now present sensitivity analysis from the prostate-cancer screening perspective. Table 5.3.7 presents sensitivity analysis data concerned with the effect of the best GA policy on the entire U.S. male population. Table 5.3.7 presents sensitivity analysis data concerned with the effect of the best GA policy on the subpopulation of U.S. males who ultimately develop prostate cancer. While we examine the effects on two different populations, the optimal policy is optimal because of its effect on one single population, the entire U.S. male population.

It is clear from Tables 5.3.7 and 5.3.7 that the discount factor, λ , is the most influential parameter. In Table 5.3.7 we see that the best GA policy always outperforms the traditional U.S. policy, in terms of the effect of screening on the entire population, but the best GA policy is better than no screening for only roughly half of the parameter values.

λ	ϵ	μ	Expected QALYs under best GA policy	Improvement over no-screening (%)	Improvement over US policy (%)
0.9924	0.05	0.01	21.580	0.042	0.009
		0.05	21.555	0.074	0.019
		0.1	21.569	-0.009	0.246
	0.145	0.01	21.568	-0.014	0.000
		0.05	21.558	-0.060	0.079
		0.1	21.570	-0.005	0.293
	0.24	0.01	21.560	-0.051	0.005
		0.05	21.561	-0.046	0.135
		0.1	21.571	0.000	0.344
1	0.05	0.01	37.778	0.186	0.077
		0.05	37.750	0.111	0.074
		0.1	37.713	0.013	0.069
	0.145	0.01	37.732	0.064	0.029
		0.05	37.706	-0.005	0.032
		0.1	37.708	0.000	0.130
	0.24	0.01	37.706	-0.005	0.034
		0.05	37.708	0.000	0.112
		0.1	37.708	0.000	0.205

Table 5.3: Sensitivity analysis on parameters of reward function in terms of the effect on the entire U.S. male population

λ	ϵ	μ	Expected QALYs under optimal policy	Improvement over no-screening (%)	Improvement over US policy (%)
0.9924	0.05	0.01	17.198	3.06	0.29
		0.05	16.925	1.42	-1.20
		0.1	16.707	0.11	-2.34
	0.145	0.01	17.038	2.10	0.87
		0.05	16.720	0.19	-0.91
		0.1	16.687	-0.01	-0.97
	0.24	0.01	16.664	-0.14	0.19
		0.05	16.664	-0.14	0.29
		0.1	16.688	0.00	0.57
1	0.05	0.01	26.640	3.63	-1.70
		0.05	26.465	2.95	-2.28
		0.1	26.250	2.11	-2.99
	0.145	0.01	26.472	0.33	-2.47
		0.05	25.793	0.33	-2.47
		0.1	25.707	0.00	-2.71
	0.24	0.01	25.809	0.40	-0.06
		0.05	25.707	0.00	-0.39
		0.1	25.707	0.00	-0.30

Table 5.4: Sensitivity analysis on parameters of reward function in terms of the effect on the subpopulation of U.S males who ultimately develop cancer

5.4 Conclusions

Table 5.4 defines the PSA-threshold screening policy that we recommend based on our research. This recommended policy was obtained by modifying the best policy discovered by the GA that is shown in Figure 5.3.4 using the base case parameter values shown in Table 4.5.1. The modifications are slight. We rounded to the nearest half-year. And we consolidated multiple identical thresholds into one threshold where the multiple thresholds occur within an extremely short period of time. After making these modifications, the recommended PSA-threshold screening policy (shown in Table 5.4) has expected QALYs equal to 37.775 ± 0.003 for a 40-year-old U.S. male; whereas, the unmodified policy shown in Figure 5.3.4 has expected QALYs equal to 37.774 ± 0.003 for a 40-year-old U.S. male. Therefore, the recommended policy, which is reasonable to implement in the clinical setting, is statistically no different from the best policy discovered by the GA with respect to expected QALYs.

Age (years)	PSA threshold (ng/mL)
42	10.0
43	10.0
50	10.0
51	10.0
55	10.0
56	10.0
57	10.0
59	9.0
62	9.0
63	9.0
64	6.5
65	6.0
66	2.5
67	2.5
67 $\frac{1}{2}$	1.0
68	0.0

Table 5.5: Recommendation for PSA-threshold screening policy

In our recommended policy it is interesting to note that the PSA threshold decreases with age, implying more intensive screening at older ages. However, screening stops at age 68. The concept of discontinuing screening at a later age is also supported by the United States Preventative Task Force, which recommends discontinuing screening at age 75. Our recommended policy

suggests screening on average every 1.7 years from age 42 to 68. This to some degree resembles the traditional U.S. policy of annual screening from the age of 50 to 75.

We now briefly discuss the most significant limitations in our study. First, the recommended policy doesn't factor in any treatment modalities, such as brachytherapy, other than prostatectomy. The original dataset, which is arguably the best currently available dataset, did not contain the information necessary to consider other treatments. Second, the population on which the dataset is based is largely caucasian, and presumably has better access to healthcare than the larger U.S. population. This is significant, particularly in light of the fact that the medical literature treats African-Americans as having a higher risk of prostate cancer than caucasians. Third, our model does not incorporate costs. While our recommended policy is optimal as measured by QALYs, we do not know the degree to which it is cost-effective.

6.1 Summary of Conclusions

For convenience, we again define the previously proposed Ross policies [39] as follows: (A) no screening; (B) threshold 4.0 ng/mL tested every 5 years from age 40 to 75; (C) threshold 4.0 ng/mL tested every 2 years from age 50 to 75; (D) annual testing with threshold 3.5 ng/mL from age 50 to 59, 4.5 ng/mL from age 60 to 69, and 6.5 ng/mL from age 70 to 75; (E) threshold 4.0 ng/mL tested annually from age 50 to 75; (F) threshold 2.5 ng/mL tested annually from age 50 to 75; (G) threshold 4.0 ng/mL tested at age 40, age 45, and every 2 years from age 50 to 75; and (H) threshold 4.0 ng/mL tested at age 40, age 45, and annually from age 50 to 75.

Policy	Expected QALYs ($\alpha = 0.05$)
A	37.708 ± 0.003
B	37.722 ± 0.003
C	37.737 ± 0.003
D	37.742 ± 0.003
E	37.749 ± 0.003
F	37.748 ± 0.003
G	37.735 ± 0.003
H	37.747 ± 0.003

Table 6.1: Comparison of the expected QALYs for the Ross policies

The recommended policy, based on the best policy discovered by the simulation-optimization method, has expected QALYs 37.775 ± 0.003 , whereas the best policy previously in use (Ross policy E, the U.S. traditional guidelines) has expected QALYs 37.749 ± 0.003 . Thus the simulation-optimization method results in an increase of approximately 0.03 QALYs on average per person. Considering the total U.S. population based on 2006 census data [1], this amounts to an estimated annual increase of 885,073 QALYs for the U.S. population. The recommended policy is repeated for convenience in Table 6.1 below.

Age (years)	PSA threshold (ng/mL)
42	10.0
43	10.0
50	10.0
51	10.0
55	10.0
56	10.0
57	10.0
59	9.0
62	9.0
63	9.0
64	6.5
65	6.0
66	2.5
67	2.5
67 $\frac{1}{2}$	1.0
68	0.0

Table 6.2: Recommendation for PSA-threshold screening policy

6.2 Implications for Research and Practice

Based on our model’s estimates of expected QALYs, there is value to PSA screening. Among the major PSA-threshold screening policies in current practice, we have shown the traditional U.S. guidelines to be superior in terms of expected QALYs. Although the traditional U.S. guidelines are superior to the other major policies in current practice, the practice of PSA-threshold screening may potentially be improved by replacing the traditional U.S. guidelines with the recommended PSA-threshold screening policy based on the best policy from our GA. If the

traditional guidelines are significantly better than no screening, then the recommended policy from our model is (most likely) significantly better than the traditional guidelines. We believe this to be true because the difference in expected QALYs between the traditional guidelines and no screening is 0.041, and the difference in expected QALYs between the recommended policy from our model and the traditional guidelines is 0.026, which is more than half the difference between the traditional guidelines and no screening.

6.3 Potential Future Research

To enhance the perspective on optimization of PSA-threshold screening, costs may be incorporated into this model, as in [47]. Providing a cost-effectiveness analysis would further inform the debate on PSA screening. Examining costs could affect the value of policies such that what is optimal with respect to QALYs is not necessarily optimal with respect to QALYs and costs.

More detailed health states would allow for more precision and accuracy in this model. States may incorporate factors such as Gleason score (which is a quantitative report on prostate biopsy), family history of prostate cancer, and other factors affecting risk of prostate cancer. Currently, this data is not available, but if it becomes available, then exploring these factors would enrich our model.

Our model is based on data from a single population of males in Olmsted County, MN. The risk of prostate cancer is not uniform among all U.S. male subpopulations. This simulation-optimization model could be extended to examine the optimal PSA-threshold screening policy for diverse populations.

Currently, our model considers only the radical prostatectomy as a treatment for prostate cancer. In reality, although the prostatectomy is the most common treatment, there are other options. This model could be modified to explore treatment strategies involving radiotherapy, brachytherapy, androgen suppressive therapy, and active surveillance. This could make our model more realistic, and could be a valuable contribution to the literature.

REFERENCES

- [1] Annual Estimates of the Population by Five-Year Age Groups and Sex for the United States: April 1, 2000 to July 1, 2006 (NC-EST2006-01). Population Division, U.S. Census Bureau, 2007.

- [2] E. Alba and M. Tomassini. Parallelism and evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 6(5):443–462, 2002.

- [3] Sigrún Andradóttir and Andrei A. Prudius. Balanced explorative and exploitative search with estimation for simulation optimization. *INFORMS Journal on Computing*, 21(2): 193–208, 2009.

- [4] G.L. Andriole et al. Mortality results from a randomized prostate-cancer screening trial. *New England Journal of Medicine*, 360(13):1310–1319, 2009.

- [5] Thomas Bäck. Selective pressure in evolutionary algorithms: A characterization of selection mechanisms. In *Proceedings of the First IEEE Conference on Evolutionary Computation*, pages 57–62. IEEE Press, 1994.

- [6] S. Baluja and R. Caruana. Removing the genetics from the standard genetic algorithm. In M.E. Kuhl, N.M. Steiger, F.B. Armstrong, and J.A. Joines, editors, *Proceedings of the Twelfth International Conference on Machine Learning*, 1995.

- [7] Peter Buchholz and Axel Thümmler. Enhancing evolutionary algorithms with statistical selection procedures for simulation optimization. In M.E. Kuhl, N.M. Steiger, F.B. Armstrong, and J.A. Joines, editors, *Proceedings of the 2005 Winter Simulation Conference*, pages 842–852, 2005.

- [8] P. Charbonneau. Genetic algorithms in astronomy and astrophysics. *The Astrophysical Journal Supplement Series*, 101(2):309–334, 1995.

- [9] Chun-Hung Chen. A lower bound for the correct subset-selection probability and its application to discrete-event system simulations. *IEEE Transactions on Automatic Control*, 41(8):1227–1231, 1996.
- [10] E. Jack Chen and Loo Hay Lee. A multi-objective selection procedure of determining a pareto set. *Computers & Operations Research*, 36:1872–1879, 2009.
- [11] Jagpreet Chhatwal, Oguzhan Alagoz, and Elizabeth S. Burnside. Optimal breast biopsy decision-making based on mammographic features and demographic factors, 2008. Working Paper.
- [12] Kenneth A. De Jong. *An analysis of the behavior of a class of genetic adaptive systems*. PhD thesis, 1975.
- [13] Kalyanmoy Deb, Samir Agrawal, Amrit Pratap, and T. Meyarivan. A Fast Elitist Non-Dominated Sorting Genetic Algorithm for Multi-Objective Optimization: NSGA-II.
- [14] Bob Djavan and Markus Margreiter. Biopsy standards for detection of prostate cancer. *World Journal of Urology*, 25:11–17, 2007.
- [15] Gerrit Draisma, Rob Boer, Suzie J. Otto, Ingrid W. van der Crujisen, Ronald A.M. Damhuis, Fritz H. Schröder, and Harry J. de Koning. Lead times and overdetection due to prostate-specific antigen screening: Estimates from the european randomized study of screening for prostate cancer. *Journal of the National Cancer Institute*, 95(12):868–878, 2003.
- [16] E.J. Dudewicz and S.R. Dalal. Allocation of observations in ranking and selection with unequal variances. *Indian Journal of Statistics*, 37(1):28–78, 1975.
- [17] Hamidreza Eskandari, Luis Rabelo, and Mansooreh Mollaghasemi. Multiobjective simulation optimization using an enhanced genetic algorithm. In *Proceedings of the 37th conference on Winter simulation*, pages 833–841. Winter Simulation Conference, 2005.

- [18] M.C. Fu. Optimization for simulation: Theory vs. practice. *INFORMS Journal on Computing*, 14(3):192–215, 2002.
- [19] M.R. Gold, D. Stevenson, and D.G. Fryback. HALYs and QALYs and DALYs, Oh My: Similarities and Differences in Summary Measures of Population Health. *Annual Review of Public Health*, 23:115–134, 2002.
- [20] David E. Goldberg. *Genetic Algorithms in Search Optimization & Machine Learning*. Addison Wesley Longman, Inc., 1989.
- [21] David Goldsman and Barry J. Nelson. Ranking, selection and multiple comparisons in computer simulation. 1994.
- [22] G.P. Haas, N.B. Delongchamps, R.F. Jones, V. Chandan, A.M. Serio, A.J. Vickers, M. Jumbelic, G. Threatte, R. Korets, H. Lilja, and G. de la Roza. Needle biopsies on autopsy prostates: Sensitivity of cancer detection based on true prevalence. *Journal of the National Cancer Institute*, 99(19):1484–1489, 2007.
- [23] N.R. Harvey and S. Marshall. The use of genetic algorithms in morphological filter design. *Signal Processing: Image Communication*, 8(1):55–71, 1996.
- [24] Randy L. Haupt and Sue Ellen Haupt. *Practical Genetic Algorithms*. Wiley, 2nd edition, 2004.
- [25] F. Herrera, M. Lozano, and A.M. Sánchez. A taxonomy for the crossover operator for real-coded genetic algorithms: an experimental study. *International Journal of Intelligent Systems*, 18:309–338, 2003.
- [26] John H. Holland. *Adaptation in natural and artificial systems : an introductory analysis with applications to biology, control, and artificial intelligence*. Ann Arbor University of Michigan Press, 2nd edition, 1975.

- [27] Koichiro Inoue, Stephen E. Chick, and Chun-Hung Chen. An empirical evaluation of several methods to select the best system. *ACM Transactions on Modeling and Computer Simulation*, 9(4):381–407, 1999.
- [28] A. Jemal, R. Siegel, E. Ward, et al. Cancer statistics, 2009. *CA: A Cancer Journal for Clinicians*, 59(4):225–249, 2009.
- [29] L.W. Koenig and A.M. Law. A procedure for selecting a subset of size m containing the l best of k independent normal populations, with applications to simulation. *Communications in Statistics - Simulation and Computation*, 14(3):719–734, 1985.
- [30] Averill M. Law. *Simulation Modeling and Analysis*. McGraw-Hill, 4th edition, 2007.
- [31] Tze-Wey Loong. Understanding sensitivity and specificity with the right side of the brain. *British Medical Journal*, 327:716–719, 2003.
- [32] Lisa M. Maillart, Julie Simmons Ivy, Scott Ransom, and Kathleen Diehl. Assessing dynamic breast cancer screening policies. *Operations Research*, 56(6):1411–1427, 2008.
- [33] Robert A.J. Matthews. The use of genetic algorithms in cryptanalysis. *Cryptologia*, 17(2): 187–201, 1993.
- [34] Charles E. Metz. ROC analysis in medical imaging: a tutorial review of the literature. *Radiological Physics and Technology*, 1(1):2–12, 2008.
- [35] Nitin Muttli and Shie-Yui Liong. Superior exploration–exploitation balance in shuffled complex evolution. *Journal of Hydraulic Engineering*, 130(12):1202–1205, 2004.
- [36] Seong Ho Park, Jin Mo Goo, and Chan-Hee Jo. Receiver Operating Characteristic (ROC) curve: practical review for radiologists. *Korean Journal of Radiology*, 5(1):11–18, 2004.

- [37] Joseph C. Presti. Prostate biopsy: how many cores are enough. *Urologic Oncology: Seminars and Original Investigations*, 21:135–140, 2003.
- [38] Y. Rinott. On two-stage selection procedures and related probability-inequalities. *Communications in Statistics – Theory and Methods*, A7(8):799–811, 1978.
- [39] K.S. Ross, H.B. Carter, J.D. Pearson, and H.A. Guess. Comparative efficiency of prostate-specific antigen screening strategies for prostate cancer detection. *Journal of the American Medical Association*, 284(11):1399–1405, 2000.
- [40] M.G. Sanda and I.D. Kaplan. A 64-year-old man with low-risk prostate cancer: review of prostate cancer treatment. *Journal of the American Medical Association*, 301(20):2141–2151, 2009.
- [41] F.H. Schröder et al. Screening and prostate-cancer mortality in a randomized european study. *Communications in Statistics - Simulation and Computation*, 3(14):719–734, 1985.
- [42] William M. Spears. *Foundations of Genetic Algorithms*, pages 221–237. 2. Morgan Kaufmann Publishers, Inc, 1993.
- [43] James R. Swisher, Sheldon H. Jacobson, and Enver Yücesan Insead. Discrete-event simulation optimization using ranking, selection, and multiple comparison procedures: A survey. *ACM Transactions on Modeling and Computer Simulation*, 13(2):134–154, 2003.
- [44] J.A. Vasconcelos, J.A. Ramírez, R.H.C. Takahashi, and R.R. Saldanha. Improvements in genetic algorithms. *IEEE Transactions on Magnetics*, 37(5):3414–3417, 2001.
- [45] Darrell Whitley and Timothy Starkweather. GENITOR II: a distributed genetic algorithm. *Journal of Experimental & Theoretical Artificial Intelligence*, 2(3):189–214, 1990.
- [46] J. Zhang, B. Denton, H. Balasubramanian, B. Inman, and N. Shah. Optimization of PSA-based screening for prostate cancer, 2009. Working Paper.

- [47] J. Zhang, B. Denton, H. Balasubramanian, N. Shah, and B. Inman. Estimating the True Value of PSA Tests for Prostate Cancer, 2010. Working Paper.
- [48] Jinghui Zhong, Xiaomin Hu, Min Gu, and Jun Zhang. Comparison of performance between different selection strategies on simple genetic algorithms. In *Proceedings of the 2005 International Conference on Computational Intelligence for Modelling, Control and Automation, and International Conference on Intelligent Agents, Web Technologies and Internet Commerce (CIMCA-IAWTIC'05)*. IEEE Press, 2005.