ANALYSIS OF BREAST CANCER SCREENING POLICIES USING PARTIALLY OBSERVABLE MARKOV DECISION PROCESSES

by

Mucahit Cevik

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Industrial and Systems Engineering)

at the

UNIVERSITY OF WISCONSIN–MADISON

2016

Date of final oral examination: 06/09/16

The dissertation is approved by the following members of the Final Oral Committee:
Oguzhan Alagoz, Associate Professor, Industrial and Systems Engineering
Jeffrey T. Linderoth, Professor, Industrial and Systems Engineering
Laura Albert McLay, Associate Professor, Industrial and Systems Engineering
Burhaneddin Sandıkçı, Associate Professor, Booth School of Business, University of Chicago
Amy Trentham-Dietz, Professor, Population Health Sciences

To my wife Merve

ACKNOWLEDGMENTS

First I must thank Professor Oguzhan Alagoz for his help and generous financial support throughout my PhD. studies. I am also grateful to Professor Burhaneddin Sandikci for his expert guidance, constant help and sincere advice. In addition, I would like to thank my dissertation committee members, Professors Laura Albert McLay, Jeffrey Linderoth and Amy Trentham-Dietz for their suggestions and insights to improve this disertation.

I am thankful to the CISNET collaborators Brian Sprague and Natasha Stout. They have always been very kind and significantly helped me expand my knowledge on the clinical aspects of my research. I would like to thank my lab mates Mehmet Ayvaci, James Codella, Mehmet Ali Ergun, Ali Hjaar, John Oruongo, and Sait Tunc for sharing the ups and downs of the PhD. life. Finally, I am grateful to my wife Merve for contributing to the every aspect of this dissertation. Without her help and constant support, this work would never had been possible.

CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

Breast cancer, the leading cause of cancer death for women, can be detected at earlier stages through screening. Therefore, several countries have implemented population-based mammography screening programs. While mammography is the gold standard for breast cancer screening, it has several drawbacks such as high rates of false positives that lead to patient anxiety and additional costs. Besides, most of the existing screening guidelines ignore individual risk factors for breast cancer, which may result in less frequent screening for high-risk women and unnecessary screenings for low-risk women. Therefore, careful design of breast cancer screening is crucial to minimize the potential harms of the screening and improving the health outcomes.

In this dissertation, we study three aspects of the breast cancer screening problem: impact of breast density and supplemental screenings, breast cancer screening in resource-restricted settings and racial disparities in breast cancer outcomes. We first analyze the impacts of breast density and supplemental tests on breast cancer screening policies. We formulate the optimal breast cancer screening problem using a discrete-time partially observable Markov decision process (POMDP) model. The state space of our model is composed of the patient's health states and the breast density states. At each decision epoch, the physician first decides whether or not the patient should undergo mammography screening, and then uses the observed mammography result to decide whether or not to follow up with supplemental screening. Our numerical study demonstrates that incorporating breast density into the design of breast cancer screening policies can significantly affect the screening recommendations. In addition, we find that incremental benefit of supplemental tests over digital mammography is rather limited; in particular, patients with higher risk of breast cancer should be recommended more frequent mammography screenings instead of supplemental tests.

Next, we investigate the optimal allocation of limited mammography resources to screen a population. We propose a constrained POMDP model that maximizes total expected quality-adjusted life years of the patients when they are allowed only a limited number of mammography screenings. We use a variable resolution grid-based approximation scheme to convert the constrained POMDP model into a mixed-integer linear program and conduct several numerical experiments using breast cancer epidemiology data. We observe that as mammography screening capacity decreases, patients in the 40-49 age group should be given the least priority with respect to screening. We further find that efficient allocation of available resources between patients with different risk levels leads to significant quality-adjusted life year gains, especially for the patients with higher breast cancer risk.

Finally, we consider race as a risk factor for breast cancer and investigate the contributing

factors leading to higher breast cancer mortality among black women. We modify the University of Wisconsin Breast Cancer Simulation model to obtain race-specific models and analyze the differences in disease natural history, treatment utilization and mammography uptake. Our findings indicate that targeted prevention and detection strategies that go beyond equalizing access to mammography may be needed to eliminate racial disparities.

1 INTRODUCTION

Breast cancer is the second most common cancer that affects the women after skin cancer. In the United States, one in eight women is expected to develop breast cancer in their lifetime and 14.5% of female cancer deaths are attributed to breast cancer [Siegel et al., 2015]. A major reason for the high number of deaths is that breast cancer is an asymptomatic disease (i.e., noticeable symptoms are typically revealed too late for successful treatment). However, numerous studies in the medical literature emphasize the importance of early detection to improve overall patient survival, and strongly support screening to help achieve this goal. In particular, American Cancer Society (ACS) reports that detecting the disease at early stages increases five-year survival rates from 27% to 98% [ACS, 2011].

There are several screening methods for breast cancer such as clinical breast examination, mammography, Magnetic Resonance Imaging (MRI) and ultrasound computed tomography. Among these methods, the most commonly used screening modality for breast cancer is mammography. On one hand, many studies report that mammography has the potential to reduce breast cancer mortality by 20%-30% through early detection of the cancer (Kerlikowske et al. [1995]). On the other hand, mammography screening has several drawbacks including exposure to radiation, false-positive test results, and risk of overdiagnosis and overtreatment. For instance, almost half of the women who are screened annually are expected to experience a false-positive mammogram within 10 years since the false-positive rate of screening mammography is around 10% [Elmore et al., 1998]. Emphasizing such issues, a recent study by the Swiss Medical Board recommends to not to introduce any new screening programs for breast cancer in Switzerland, and furthermore, to phase out existing programs [Biller-Andorno and Jüni, 2014].

Screening exams are administered to detect breast cancer in women who have no apparent symptoms. If the screening exam leads to suspicious results, the physician usually follows up with a diagnostic test to check the tissue. Diagnostic mammogram is the most commonly used diagnostic test to evaluate abnormalities detected on a screening exam. It differs from

a screening mammography in that additional views of the breast are taken, and therefore a more detailed x-ray of the breast tissue is provided. In addition, MRI and ultrasound can also be used as diagnostic tests. In this dissertation, we do not consider the post-screening diagnostic decisions and focus on the breast cancer screening problem. More information about optimization approaches in breast cancer diagnosis can be found elsewhere [Chhatwal et al., 2010, Ayvaci et al., 2012].

Risk factors associated with the breast cancer further complicates assessment of the benefits and harms of the screening. For instance, some studies note that breast cancer is known to be more aggressive for younger women and suggest more frequent screening at earlier ages [Jayasinghe et al., 2005]. However, other studies suggest more frequent screening for older women considering that nearly 50% of the newly diagnosed cancers occur in women older than 60 and mammography screening have higher accuracy for older women [ACS, 2011, Kerlikowske et al., 2000]. Reflecting the aforementioned controversies, there exist significant inconsistencies in mammography screening guidelines in the world. Several countries have different guidelines in terms of starting age and ending age, as well as the frequency of the screening. For instance, patients are recommended biennial mammography screening between ages 45 and 69 in Spain, and triennial mammography screening between ages 50 and 70 in the United Kingdom [Klabunde et al., 2007]. Similar inconsistencies in mammography screening guidelines may exist even within the same country. For example, in the U.S., while U.S. Preventive Services Task Force (USPSTF) recommends biennial mammography screening between ages 50 and 74, American College of Radiology (ACR) recommends annual mammography screening, starting at age 40 [USPSTF, 2009, Lee et al., 2010]. Table 1.1 summarizes different screening guidelines by various medical organizations in the US and selected countries that have population-based screening programs.

Screening a population of individuals may be beneficial in detecting diseases at an early stage. However, majority of women would not receive any benefits from screening, since only a small portion of the screened population would develop breast cancer throughout their lifetime.

Institution/Country	Start Age	End Age	Screening intervals in years
American Cancer Society	40	Not specified	1
American Medical Association	40	Not specified	1
American College of Radiology	40	Not specified	1
National Cancer Institute	40	Not specified	1 - 2
U.S. Preventive services Task Force	50	74	2
American College of Preventive Medicine	50	Not specified	1 - 2
American Academy of Family Physicians	50	Not specified	1 - 2
Canada, Italy, Japan	50	69	2
France, Netherlands	50	74	2
Spain	45	69	2
United Kingdom	50	70	3

Table 1.1: Recommended mammography screening policies by various institutions in the United States and other countries with organized population-based cancer-screening programs (adopted from Ayer et al. [2012])

As such, the main weakness of population-based screening guidelines remains to be ignoring an individual woman's breast cancer risk factors other than age. Gail and Rimmer [1998] note that each woman has different personal characteristics that affect breast cancer risk; these factors include family history, breast density, body mass index, alcohol consumption, parity, extent of breastfeeding and ages at menarche, menopause and first birth. For instance, a woman with a family history of breast cancer is two or three times more at risk of getting breast cancer [Gilbar, 1998]. As a result, screening guidelines tailored to individual risk factors have potential to increase life-savings in high-risk women, while preventing unnecessary mammograms and the resulting harms in low-risk women. Moreover, rising screening costs and mounting economic pressure further necessitates optimizing breast cancer screening recommendations. One study reports that the combined cost of mammography screening along with the work-up costs of positive findings is estimated to be $3 billion to $5 billion dollars in the United States [Burnside et al., 2001].

The main focus of this dissertation is to deal with the question of how to improve breast

cancer screening recommendations based on individual risk factors. In this regard, we study three problems faced by policy makers: eliminating racial disparities in breast cancer outcomes, improving screening recommendations using breast density information and supplemental screening methods, and analyzing the impact of limited resources on breast cancer screening recommendations. In the remainder of this chapter, we first discuss the role of breast density as a risk factor and potential of supplemental screening methods on improving health outcomes in Section 1.1. Then, in Section 1.2, we describe the breast cancer screening problem in resource-restricted settings. We summarize the role of racial disparities on breast cancer incidence and mortality in Section 1.3 and conclude the chapter in Section 1.4 with an overview of the thesis and contributions.

1.1 Role of Breast Density and Supplemental Screening Methods on Breast Cancer Screening Policies

Mammography is the most commonly used screening method for breast cancer. A mammography is an x-ray image of the breast that makes it possible to detect tumors that cannot be felt. The image can be captured on film or stored directly onto a computer (digital). As the digital mammography is viewed on a computer, it has the added benefit of lightening or enlarging the images to better examine the breast. Both film and digital mammographies are known to be less accurate for women with dense breasts for whom cancer may be masked by dense breast tissue [Lee et al., 2014]. Digital mammography largely replaced the film mammography in many U.S. clinics particularly because of its improved sensitivity for women with dense breasts [Pisano et al., 2005]. However, as the specificity of digital mammography is modestly lower than film mammography, overall health benefits of digital mammography remain unclear [Kerlikowske et al., 2011, Stout et al., 2014].

Because of the imperfect nature of mammography, patients may generally seek supplemental

screening methods, such as Magnetic Resonance Imaging (MRI) and ultrasound, to achieve better detection accuracy [Berg et al., 2012]. However, there are ongoing debates about the cost-effectiveness of the use of supplemental tests in breast cancer screening. While there exists a group of physicians that frequently refer patients for supplemental screenings, others use these tests more conservatively, which is most likely due to a lack of clear guidelines. Moreover, additional costs incurred due to supplemental tests may also affect the physician's recommendation.

Recent medical literature identifies a patient's estimated breast cancer risk along with the patient's breast density as key factors influencing the decision to use supplemental tests [Berg et al., 2012, Lee et al., 2014, Sprague et al., 2014]. Regardless of their breast density, patients with high breast cancer risk may be recommended to undergo MRI screening because of its high sensitivity. However, due to its low-availability, high cost, and low-specificity, use of MRI for the general population is limited [Berg, 2014]. On the other hand, ultrasound is widely available and has relatively low cost compared to MRI. As a result, it is often the preferred supplemental screening test for women with high breast cancer risk. Despite their potential benefits especially for high-risk patients, the effect of supplemental tests on long-term health outcomes is still debated. In a recent study, Sprague et al. [2014] conclude that supplemental ultrasound screening for women with dense breasts substantially increases costs while producing relatively small benefits. Similarly, Melnikow et al. [2016] conclude that supplemental screening of women with dense breasts would increase the false-positive results considerably and more research is needed to understand the effects of supplemental screenings on breast cancer outcomes.

Along with several other factors such as family history and obesity, breast density is considered to be a strong risk factor for breast cancer [Wang et al., 2014]. As a result of its significance, some states in the U.S. enacted laws that require patients to be informed about their breast density. Besides, breast density has a dynamic nature — that is, it changes as the patient gets older. For instance, Sprague et al. [2014] report that while 61.4% of women younger than 50 have dense breasts, this ratio drops to 40.8% for women older than 50. Therefore, optimization of

breast cancer screening policies is further complicated by the incorporation of breast density information to the decision process.

1.2 Role of Limited Resources on Breast Cancer Screening Policies

Cost of mammography and lack of resources in terms of number of diagnostic machines and the number of trained workforce to interpret these mammograms (i.e., radiologists) limit the widespread use of mammography for breast cancer screening in many developing countries [BahadurSingh et al., 2014]. In fact, many countries face capacity issues even when the cost of screening is of no concern. For example, according to the 2013 census of China, there are 117 million Chinese women between the ages 50-69 [China Statistical Yearbook, 2015]. Implementing biennial screening in this age group only (which is a less aggressive screening policy than most policies adopted by developed countries) would require adding more than 30,000 radiologists to the overall workforce, assuming a radiologist can interpret approximately 2000 mammographies per year [Smith-Bindman et al., 2008], which would necessitate increasing the number of radiologists in China by nearly 50% [IRQN, 2015]. Similar issues exist in other developing countries as well. For instance, Reddy et al. [2012] note that a nationwide mammographic screening program in India is infeasible primarily due to the poor infrastructure and the costs involved in screening.

Although several developed countries implement population-based mammography screening guidelines, the majority of developing and underdeveloped countries do not have mammography screening programs. In fact, WHO recommends mammography screening only for resource-rich countries [WHO, 2015]. However, even in resource-rich countries such as in the U.S., there is a growing concern about the cost of mammography screening. For example, California Department of Public Health, the state-run program providing free mammography screenings to uninsured women, raised the eligibility age for breast cancer screening services from 40 to 50 for women in the Every Woman Counts program effective from January 1, 2011 [Schneider, 2010]. The primary reason for dropping coverage for these women was cited as "the

unprecedented fiscal challenge" as a result of increasing demand for breast cancer screening and declining revenues. Susan G. Komen Foundation also reports that there have been significant budget cuts in mammography screening funds for low-income women in several states including California, Colorado, Michigan, New York, Ohio, Pennsylvania, and Washington [Susan G. Komen Foundation, 2011].

Another important factor to consider in allocating limited mammography screening resources is the variation among patients with respect to breast cancer risk. For example, patients with family history are twice or three times more likely to develop breast cancer compared to ones with no family history of breast cancer [Gilbar, 1998]. Therefore, high-risk patients such as those with a family history are likely to benefit more from mammography screening. Considering this variation in risk may be particularly important when resources are limited. For instance, given a very limited budget, a policy maker may prioritize screening patients with family history over those without a family history.

1.3 Role of Racial Disparities on Breast Cancer Incidence and Mortality

Racial disparities between African-American and European-American women in the U.S. regarding breast cancer incidence and mortality have long been a public health concern. Despite lower incidence rates, breast cancer mortality rates have been higher among black women for the past 30 years and disparity in the mortality between black and white women continues to grow [Ghafoor et al., 2002, Smigal et al., 2006]. For example, according to Surveillance Epidemiology and End Results (SEER), mortality rate for black women was 31 per 100,000 women in 2008, whereas it was 22 for whites, which points to 40% difference in mortality rates [Howlader et al., 2011].

Multiple factors are associated with poorer survival among black women, including histori-

cally lower rates of mammography screening, decreased likelihood of receiving stage appropriate treatment, more comorbid conditions and higher rates of obesity [Vastag, 2003, Newman, 2005, Chlebowski et al., 2005, Amend et al., 2006]. In addition, several studies suggest that black women are usually diagnosed with breast cancer at a more advanced stage and are more likely to have tumors with worse prognosis than whites, which suggest that variation in tumor natural history might be a significant contributing factor for racial disparities [Wojcik et al., 1998]. In this regard, understanding the racial variation in breast cancer biology, racial differences in actual treatment utilization, and the differences in screening uptake is important to refine optimal strategies for eliminating racial disparities in breast cancer outcomes.

1.4 Thesis Overview

We close the introduction with a detailed overview of the thesis, which contributes to improving breast cancer screening policies.

In Chapter 3, we evaluate the benefits of supplemental screening tests through incorporating breast density information and recommend screening policies that maximize an individual patient's total quality-adjusted life years (QALYs). We model the optimal breast cancer screening problem as a discrete-time partially observable Markov decision process (POMDP). We estimate the parameters of our model from University of Wisconsin Breast Cancer Simulation Model (UWBCS), which uses breast cancer epidemiology data. Our model provides insights to eliminate unnecessary supplemental tests and help physicians make more personalized screening recommendations.

In Chapter 4, we investigate the breast cancer screening problem in a resource-constrained setting, which is often the case for developing countries but it can also include screening programs in developed countries such as the Every Woman Counts program in California. We seek answers to many important health policy questions such as: How should limited mammography screening resources be allocated for patients in different age groups? Given a limited capacity, should

we screen high-risk groups aggressively and ignore low-risk population or screen high-risk patients less aggressively and allocate some resources to the low-risk groups? What would be the impact of increasing screening capacity on the overall health outcomes? In order to address these questions, we propose a constrained POMDP modeling framework with the objective of maximizing the total expected QALYs of the patients. We analyze the impact of limited resources at both population and individual levels. We first model an individual patient's breast cancer screening problem considering that only a limited number of mammography screenings is available per patient. Then, we extend our model to consider different risk groups and obtain a resource allocation scheme between patients with different risk levels in a given population. To our knowledge, even unconstrained POMDPs are notorious for being computationally intractable and there is no known efficient solution algorithm to solve constrained POMDPs to optimality. In that regard, we propose a grid-based approximation scheme to solve the constrained POMDP model using a mixed-integer linear program. We estimate the parameters of our model from a validated simulation model that uses breast cancer epidemiology data and compare the performance of our proposed policies to those of various population-based screening policies. Although we primarily use U.S.-based data sources, our approach and general policy insights apply to other resource-limited settings as well.

In Chapter 5, we investigate the underlying factors causing the racial disparities in breast cancer incidence and mortality. We classify the contributing factors into three main categories: natural history of the disease, screening uptake and treatment utilization. We modify UWBCS, a previously validated microsimulation model for breast cancer, to obtain race-specific models. We use these race-specific models to assess the contributions of each factor on the breast cancer mortality difference between black and white women. Obtaining the race-specific models require a lengthy calibration procedure, where we adjust some of the natural history parameters of the UWBCS to replicate U.S.-observed breast cancer incidence and mortality rates for black and white women. The natural history parameters (e.g., mean tumor growth rate and percentage of benign

tumors) usually cannot be informed by the existing data sources and their values are largely unknown. A common approach for calibration is to generate a set of natural history parameter combinations and evaluate these parameter combinations via simulation to obtain the best fitting parameter values. UWBCS has ten natural history parameters and considering that each parameter takes five values, the number of parameter combinations would be approximately 9.7 millions (5^{10}). In this study, we propose a novel method using machine learning approaches to speed up the calibration procedure. We further discuss that our calibration method is applicable to a generic simulation model that requires evaluation of large number of parameter combinations for the calibration.

2 LITERATURE REVIEW

In this chapter, we review the related literature. Section 2.1 introduces the POMDP models and their applications in healthcare. Section 2.2 summarizes the analytical models in cancer screening with an emphasis on breast cancer screening. Finally, Section 2.3 describes the studies on racial disparities in breast cancer outcomes.

2.1 Partially Observable Markov Decision Processes

Partially observable Markov decision process (POMDP) models relax the completely observable system state assumption in Markov decision processes (MDP). As such, probability of being in a particular state is taken as the basis for choosing an action. POMDPs are useful for problems for which actions can help modifying system trajectory in a desired manner and gathering information for future decisions.

POMDPs have many application areas including inventory control [Treharne and Sox, 2002], inspection of structural units (e.g., roads and bridges) [Ellis et al., 1995], search for moving objects [Eagle, 1984], and machine maintenance and replacement [Eckles, 1968, Maillart, 2006]. Furthermore, POMDPs have been frequently used in medical decision making problems such as medical diagnosis and treatment [Hu et al., 1996, Hauskrecht and Fraser, 2000], healthcare system analysis [Smallwood et al., 1971] and cancer screening [Ayer et al., 2012, Zhang et al., 2012]. Monahan [1982] provides an overview of POMDPs and some of the potential applications.

POMDP models are notoriously difficult to solve optimally. The main difficulty in solving the POMDP models stems from the fact that the belief space is a continuum. As such, we have infinitely many equations with infinite number of variables in the optimality equations. A key observation to improve solvability of the POMDP models is that the POMDP value functions have piecewise linear convex representations. More specifically, Smallwood and Sondik [1973] show that POMDP value functions ($V(\pi)$ for belief state π) can be represented as $V(\pi) = \alpha\pi$

and propose the first exact solution algorithm that aims to construct minimal α-vector set. Later, several other studies extend Sondik's algorithm to solve POMDPs optimally, [Monahan, 1982, Cheng, 1988, Cassandra et al., 1997]. It is important to note that the majority of these optimal solution algorithms rely on piecewise linear representation of the value functions.

Dimensionality and precision issues render most POMDP models intractable by exact solution methods. Therefore, several other studies propose approximation algorithms. In particular, finite grid approximation to the uncountable state space is a frequently used approximation technique for POMDPs. Lovejoy [1991b] uses a fixed grid method to generate upper and lower bounds for the optimal value function. Sondik and Mendelssohn [1979] uses a variable grid method that allow the grid points used in the approximation algorithm to vary from one iteration to next. We refer the readers to Lovejoy [1991a] for a detailed survey of available algorithms.

2.2 Analytical Models in Cancer Screening

Operations research tools are well utilized in the literature to determine efficient cancer screening schedules for a population or an individual patient. A comprehensive survey on modeling and analysis of cancer screening problems can be found in Alagoz et al. [2011] and Ivy [2009]. We restrict our review to the papers addressing stochastic models in cancer screening.

Most studies in the literature propose simulation models to evaluate whether a particular population-based screening strategy such as annual mammography screening between ages 40-74 is cost-effective or not in a resource-rich country. For example, National Cancer Institute's Cancer Intervention and Surveillance Modelling Network (CISNET) consortium involves six breast cancer models that use simulation/statistical tools to make recommendations to policy makers in the U.S. [NCI, 2013].

In addition to the simulation models, several studies use Markovian models to investigate cancer screening decisions. Among such models, partially observable Markov chains (POMC) have gained attention for modeling cancer screening problems as they can capture age-based

dynamics of cancer screening and the effects of imperfect test results. POMCs have been used by Maillart et al. [2008] and Madadi et al. [2015] to compare various screening policies for breast cancer. Also, Li et al. [2014] use a POMC model to evaluate a broad range of colonoscopy screening schedules.

Some other studies develop Markov decision process (MDP) models to optimize screening decisions. For instance, Chhatwal et al. [2010] and Ayvaci et al. [2012] model breast cancer biopsy decision-making problem as MDPs. In particular, Ayvaci et al. [2012] develop a constrained MDP model to identify optimal post-mammography diagnostic decisions for an individual patient under budgetary restrictions. As they do not aim to optimize the screening decisions, they disregard the imperfect test results and assume that the model states are fully observable. They conduct an extensive numerical study and find that diagnostic-decision-making thresholds increase significantly as the budget gets tighter.

As POMDPs generalize MDPs by allowing incomplete information about the state of the process, which is a natural situation in most cancer screening problems, they have been frequently used in several studies. In particular, studies that consider cancer screening problem frequently aim to identify optimal screening decisions from an individual patient's perspective to maximize the health outcomes by balancing the benefits of screening via early detection and harms caused by false-positive test results. For example, Erenay et al. [2014] develop a POMDP model to identify the optimal screening schedules for colorectal cancer for both individual men and women, and Zhang et al. [2012] propose a POMDP model to optimize prostate cancer biopsy referral decisions for individual men. In this body of work, the closest studies to our work are those by Ayer et al. [2012, 2016]. Ayer et al. [2012] propose a POMDP model to individualize the breast cancer screening problem, and Ayer et al. [2016] extend this work to incorporate the patient's adherence behavior to the screening recommendations. However, Ayer et al. [2012, 2016] do not consider the impact of breast density and availability of supplemental screening methods on the personalized breast cancer screening policies.

In the Operations Research/Management Science literature, a limited number of studies have investigated the impact of resource restrictions in cancer screening. Güneş et al. [2015] propose a compartmental model for the allocation of limited colonoscopy resources between screening and diagnostic services, assuming that patients are recommended periodic screening with colonoscopy every ten years. Lee et al. [2014] use reinforcement learning to design policies for hepatocellular carcinoma screening where k out of n patients are chosen for screening at each period. Our main difference from these studies is that we consider both imperfect test results and resource constraints simultaneously in our modeling framework. As such, this leads to a different modeling and solution approach, which can be used for determining the optimal resource allocation between risk-groups, and optimal screening strategy within each group. To this end, our proposed approach is unique in that it not only provides a framework to optimize the breast cancer screening problem under resource constraints, but also provides a modeling framework to optimize resource allocation decisions in other screening problems, where disease status is only partially observable.

2.3 Racial Disparities

Racial disparities in breast cancer outcomes have been the subject of many studies over the years. Majority of these studies carry out retrospective analysis and use available data sources to investigate potential contributing factors to racial disparities. Smith-Bindman et al. [2006], Jazieh and Buncher [2002], Elmore et al. [2005] observe that mammography screening uptake and mammography effectiveness due to differences in breast density distributions might vary significantly between black and white women. Morris et al. [2007], Chlebowski et al. [2005], Banerjee et al. [2007], Carey et al. [2006], Bauer et al. [2007], Wojcik et al. [1998] find that black women are more likely to be diagnosed with breast cancer at a later stage and are more likely to have tumors with poorer prognosis such as Estrogen Receptor (ER) negativity, poorer differentiation and greater lymph node involvement. Furthermore, Elmore et al. [2005], Fedewa et al.

[2010], Hershman et al. [2005, 2009] suggest that black women experience more delays between cancer detection and treatment initiation and they are also more likely to terminate treatment prematurely, and Shavers and Brown [2002], Joslyn [2002], Li et al. [2003] conclude that black women are less likely to undergo required surgeries and follow-up radiation therapy after the surgery.

There are only a few studies in the literature that use analytical models to investigate racial disparities in breast cancer outcomes. In a recent study, van Ravesteyn et al. [2011] report that the higher breast cancer mortality rates for black women could be attributed to differences in natural history parameters (26-44%), use of adjuvant therapy (11-19%) and uptake of mammography screening (7-8%) and they conclude that 38% to 46% of the difference cannot be explained by their models. Later, Batina et al. [2013] evaluate the contribution of racial differences in tumor natural history to observed disparities in breast cancer incidence. Their findings indicate that mean tumor growth rate is found to be 63.6% higher, and percentage of highly aggressive tumors are 2.2 times greater for black women compared to the white women.

3 IMPACT OF BREAST DENSITY AND SUPPLEMENTAL SCREENING METHODS ON BREAST CANCER SCREENING POLICIES

3.1 Problem Definition

Breast cancer is a disease in which cancer cells form in the tissues of the breast. Since there are usually no physical symptoms of breast cancer, patients usually undergo mammography screening to detect the disease. There are several critical decisions involved in the design of breast cancer screening policies. These decisions are complicated by the changes in the breast density of the patient, probabilistic progression of breast cancer, imperfect test results, and quality-of-life considerations associated with screening decisions. In addition, there are no clear guidelines about the use of supplemental screening tests in timely detection of breast cancer. There have been recent advances in using mathematical optimization models for specific aspects for the breast cancer screening problem. However, optimal policies that simultaneously consider a patient's breast density and the availability of supplemental tests are not yet well understood. Our goal in this paper is to use optimization modeling to better understand the impact of these factors on breast cancer screening recommendations.

Figure 3.1 provides an overview of the typical screening process that is commonly used in the U.S. medical practice. The screening process consists of making decisions at regular intervals — typically every 6 months, annually, biennally, or triannually. There is a nonzero state-dependent likelihood of patient death during each screening interval, causes of which include breast cancer as well as others. In each screening interval t, the physician first decides whether the patient, with some breast cancer risk assessment $\boldsymbol{\pi}_t$, should receive mammography screening (M) in the current period or wait (W) until the next decision epoch. If the wait action is chosen, the patient may self-diagnose herself (i.e., through feeling a suspicious lump on her breast or via clinical breast exam) and, using the outcome of her self-diagnosis, she constructs a posterior $\boldsymbol{\pi}_{t+1}$ and faces the same mammography screening decision in the new decision epoch $t+1$, provided that

she survives.

If, on the other hand, mammography screening action is chosen, it can result in a positive $(+)$ finding, which indicates the existence of some abnormal suspicious tissue that may or may not be an actual cancer, or a negative $(-)$ finding, which indicates a lack of any evidence for potentially cancerous tissue. In addition to serving as a signal for the patient's cancer status, a mammogram also reveals the breast density of the patient, which is otherwise unknown to the physician. MRI and ultrasound can be used along with mammography screening to improve the detection accuracy. However, stand-alone usage of these tests are not recommended as they are associated with large numbers of false-negative and false-positive results [Berg et al., 2012]. Furthermore, note that MRI and ultrasound may also be used as diagnostic tests after a positive mammography result, but we assume that positive test results are immediately followed by a perfect diagnostic test such as biopsy, which correctly identifies the cancer status of the patient, and therefore, we only focus on using such supplemental tests after a negative mammogram. If the diagnostic test (e.g., biopsy) confirms the positive screening test result, the patient is assumed to immediately start cancer treatment, concluding the screening decision process. On the contrary, if the diagnostic test nullifies the positive screening test result, the patient restarts the screening decision process in the next decision epoch $t+1$ with posterior $\boldsymbol{\pi}_{t+1}$, provided that she does not die of other causes.

A negative mammography result is used in two major ways. It reveals the true current breast density of the patient and helps the physician update her or his assessment of the patient's risk of breast cancer before taking a recourse action in the form of recommending a supplemental screening or doing nothing (N) until the next decision epoch. It is important to note that the timing t' of the recourse actions is shortly after the initial mammography screening decision when compared to the full duration of the decision epoch. This distinction will be critical in updating the risk assessments. If a supplemental screening test is taken, the patient follows a similar pathway to that of the initial mammography decision except that a negative supplemental test

result is not followed by any other decisions, and therefore, the patient faces the mammography decision again with an updated posterior $\boldsymbol{\pi}_{t+1}$ in the next decision epoch $t+1$, provided that she survives.

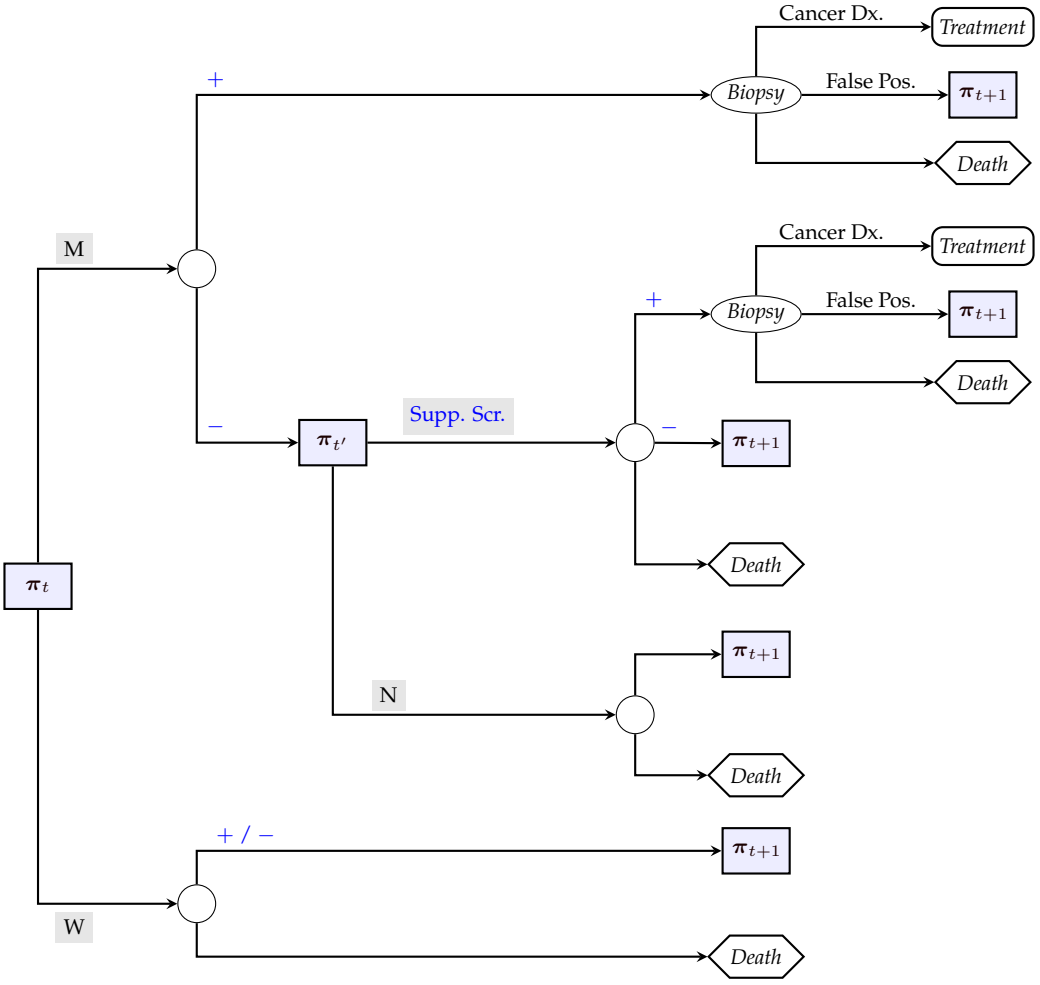

Figure 3.1: Breast cancer screening decision-making diagram

3.2 The POMDP Model

We formulate a discrete-time finite-horizon POMDP model to optimize breast cancer screening recommendations of a physician. Screening decisions are made at discrete points t over a finite time horizon $\mathcal{T} < \infty$. We define the core state space of the process as $\mathcal{S} = \bar{\mathcal{S}} \cup \{\nabla\}$, where ∇ is an absorbing state that represents death, and $\bar{\mathcal{S}} = \{d, h \mid d \in \mathcal{D}, h \in \mathcal{H}\}$. $\mathcal{D}$ is discrete set of breast-density related states, and we take $\mathcal{H} = \Delta \cup \mathcal{C}$, where Δ denotes the cancer-free state and $\mathcal{C}$ is a discrete set of breast-cancer related states. Note that health states are not directly observable, whereas breast density states are only revealed when the patient undergoes a screening exam. As a result, the physician has incomplete knowledge about patient's state $(d_t, h_t) \in \bar{\mathcal{S}}$ at any decision epoch t.

Since the decision maker cannot fully observe the core state $s \in \bar{\mathcal{S}}$, she constructs a belief $\pi(s)$, which corresponds to the probability of being in s. That is, a belief state $\boldsymbol{\pi}$ is a probabilistic construct about the physician's assessment of the patient's health and breast density. The set of all belief states is given by the simplex $\Pi(\bar{\mathcal{S}}) = \left\{\boldsymbol{\pi} \in \Re^{|\bar{\mathcal{S}}|} : \ \sum_{s\in\bar{\mathcal{S}}} \pi(s) = 1, \ \pi(s) \geq 0, \ \forall s \in \bar{\mathcal{S}}\right\}$, and $\boldsymbol{\pi} \in \Pi(\bar{\mathcal{S}})$ is updated in Bayesian manner as actions are taken and observations are made over time.

At the beginning of any decision epoch, the physician recommends either watchful waiting (W) or mammography screening (M), and he may follow-up the mammogram with a supplemental screening exam ($a \in \mathcal{A}^{Sup}$). The physician may also choose to recommend doing nothing (N) as an alternative to supplemental screening (see Figure 3.1 for an illustration). For convenience, we refer to the set of all possible actions as $\mathcal{A} = \{W, M, N\} \cup \mathcal{A}^{Sup}$ and the set of screening actions as $\mathcal{A}^{Scr} = \{M\} \cup \mathcal{A}^{Sup}$.

Taking action $a \in \mathcal{A} \setminus \{N\}$ leads to an observation that signals patient's health and breast density states. In particular, the physician makes a positive (ξ^+) or a negative (ξ^-) observation about the patient's health state according to the sensitivity and specificity of the action. If the patient takes the wait action, she does not make any observations about her breast density state

(η_{no}). Alternatively, if the patient undergoes screening, then her breast density is revealed and she makes the observation η_i at breast density state $d_i \in \mathcal{D}$. We represent the observations for the core states as $o = (\eta, \xi) \in \Omega_a$, where $\Omega_W = \{(\eta_{no}, \xi) : \xi \in \{\xi^-, \xi^+\}\}$ and $\Omega_a = \{(\eta, \xi) : \eta \in \{\eta_0, \ldots, \eta_{|\mathcal{D}|}\}, \xi \in \{\xi^-, \xi^+\}\}$ for all $a \in \mathcal{A} \setminus \{W, N\}$.

A screening test can detect in situ or invasive cancer with probabilities equal to its sensitivity for the stage of the tumor and breast density of the patient, denoted as $\tau_a^{(d,h)}$, for $a \in \mathcal{A}^{Scr}$. Similarly, patient can make a self-detection with probability $\tau_W^{(d,h)}$. If the patient does not have cancer, each of these actions can lead to a false-positive diagnosis with probability $1 - \nu_a^d$, where ν_a^d is the specificity of the corresponding action. We use $f_t^a(\eta, \xi | d, h)$ to denote the observation probabilities for the core states at decision epoch t when the patient takes the action $a \in \mathcal{A}$. A table form representation of the $f_t^a(\eta, \xi \mid d, h)$ values is given in Appendix A.2.

Transition probabilities $p_t^a(d', h' | d, h)$ denote the probability that a patient will be in state (d', h') in decision epoch $t + 1$, given that she is in state (d, h) and takes action $a \in \mathcal{A}$. Breast density is a significant risk factor for breast cancer (see Appendix A.1). Therefore, estimated transition probabilities reflects the effects of breast density on tumor onset and progression.

The immediate rewards in our model correspond to the QALYs accrued by a patient at a given decision epoch. We use the half-cycle correction method to calculate the immediate rewards. Let $\omega_t(s)$ represent the immediate reward for occupying state $s \in \bar{\mathcal{S}}$ at decision epoch t. That is,

$$\omega_t(s) = \lambda(s) \times \big[1 \times P(\text{alive in } t+1 | \text{current state is } s) + 0.5 \times P(\text{dead in } t+1 | \text{current state is } s)\big],$$

where $\lambda(s)$ is disutility multiplier for occupying state $s \in \bar{\mathcal{S}}$. Then, immediate reward at belief state $\boldsymbol{\pi} \in \Pi(\bar{\mathcal{S}})$ is calculated as follows:

$$r_t(\boldsymbol{\pi}) = \sum_{s \in \bar{\mathcal{S}}} \pi(s) \omega_t(s).$$

If the patient is diagnosed with cancer at the decision epoch t, she accrues a lump-sum reward

$\psi_t(s)$, and quits the decision process. Besides, no action is taken at the final decision epoch and the patient accrues a terminal reward $\omega_{\mathcal{T}}(s)$ in state $s \in \mathcal{S}$.

We also account for quality of life reductions due to disutilities associated with screening actions. Each screening test has disutilities that depend on the outcome of the test. Let $u(a)$ represent the the disutility of taking action $a \in \mathcal{A}^{Scr}$. Also, let μ^{TP} and μ^{FP} represent the disutilities of true-positive (TP) and false-positive (FP) screening results, respectively. We obtain the disutility value, $\gamma(s, a, o)$, associated with the screening action $a \in \mathcal{A}^{Scr}$ for each state $s = (d, h)$ and observation $o \in \Omega_a$ as follows:

$$\gamma(s=(d,h),a,o) = \begin{cases} u(a), & \text{if } o=(\eta,\xi^-) \in \Omega_a, \\ u(a)+\mu^{FP}, & \text{if } h=\Delta,\ o=(\eta,\xi^+) \in \Omega_a, \\ u(a)+\mu^{TP}, & \text{if } h\neq\Delta,\ o=(\eta,\xi^+) \in \Omega_a. \end{cases}$$

Then, expected disutility of taking a screening action $a \in \mathcal{A}^{Scr}$ at the belief state $\pi \in \Pi(\bar{\mathcal{S}})$ is

$$\gamma(\boldsymbol{\pi}, a) = \sum_{s\in\bar{\mathcal{S}}} \sum_{o\in\Omega_a} \pi(s) f^a(o|s) \gamma(s, a, o).$$

Belief state, $\boldsymbol{\pi}$, of the patient changes as the patient takes actions and makes observations about her health and breast density. Let $\pi'(s)$ represent the probability of occupying state $s \in \bar{\mathcal{S}}$ at the subsequent decision epoch when the current belief state of the patient is $\boldsymbol{\pi}$, action is $a \in \mathcal{A}$,

and observation is $o \in \Omega_a$. That is,

$$\pi'(s') = \begin{cases} \dfrac{\sum_{s\in\bar{\mathcal{S}}}\ \pi(s)\cdot f_t^a(o\mid s)\cdot p_t^a(s'|s)}{\sum_{s\in\bar{\mathcal{S}}}\ \pi(s)\cdot f_t^a(o\mid s)\cdot \mathcal{P}_t^a(s)} & \text{if } a=W, o=(\eta,\xi)\in\Omega_W, \\ & \text{or } a\in\mathcal{A}^{Sup},\ o=(\eta,\xi^-)\in\Omega_a, \\ \dfrac{\pi(s')\cdot f_t^M(o\mid s')}{\sum_{s\in\bar{\mathcal{S}}}\ \pi(s)\cdot f_t^M(o\mid s)}, & \text{if } a=M,\ o=(\eta,\xi^-)\in\Omega_M, \\ p_t^a(s'=(d',h')|s=(d,\Delta)) & \text{if } a\in\mathcal{A}^{Scr},\ o=(\eta,\xi^+)\in\Omega_a, \\ \dfrac{\sum_{s\in\bar{\mathcal{S}}}\ \pi(s)\cdot p_t^N(s'\mid s)}{\mathcal{P}_t^N(s)} & \text{if } a=N. \end{cases} \tag{3.2a–3.2d}$$

where $\mathcal{P}_t^a(d,h)=\sum_{(d',h')\in\bar{\mathcal{S}}}\ p_t^a(d',h'|d,h)$ is the probability of patient being alive at the decision epoch $t+1$.

The patient does not advance to the decision epoch $t+1$ if she undergoes mammography screening and receives a negative test result (ξ^-). Instead, she makes a transition to the immediate decision epoch t'. Therefore, we use the formula (3.2b) to update π for this case. On the other hand, if the patient receives a positive test result (ξ^+) from any screening action, her true health state is revealed. If the test is true-positive, then we do not need to update the patient's belief state as the patient quits the decision process and starts receiving treatment. If the test is false-positive, then we use the natural history progression of the disease to update the belief state as in (3.2c). Finally, if the patient does not receive a supplemental screening test, she does not make another observation in this period, and her belief state is updated according to formula (3.2d).

Our model aims to determine the breast cancer screening decisions that maximize expected QALYs over the patient's lifetime. As we allow sequential screenings in our model, we separate the value functions into two stages. We define the optimal value functions $V_t^W(\boldsymbol{\pi})$ and $V_t^M(\boldsymbol{\pi})$ as the maximum expected QALYs from year t through $\mathcal{T}$ for the first stage wait and mammography actions, respectively. Second stage value functions are only relevant if the patient takes the mammography action and makes a negative observation about her health state ($i.e., o=\xi^-$).

We use $V^N_{t,\eta}(\boldsymbol{\pi}),$ and $V^a_{t,\eta}(\boldsymbol{\pi})$ to represent the value functions for the second stage do-nothing and supplemental screening action $a \in \mathcal{A}^{Sup}$, respectively. The optimality equation for the wait action is as follows:

$$V^W_t(\boldsymbol{\pi}) = r_t(\boldsymbol{\pi}) + \sum_{(d,h)\in\bar{S}} \pi(d,h) \sum_{(\eta,\xi)\in\Omega_W} f^W_t(\eta,\xi|d,h)\mathcal{P}^W_t(d,h)V^*_{t+1}(\pi'). \tag{3.3}$$

Similarly, we obtain the optimal value function for the mammography action as follows:

$$V^M_t(\boldsymbol{\pi}) = -\gamma(\boldsymbol{\pi}, M) + \sum_{(d,\Delta)\in\bar{\mathcal{S}}} \pi(d,\Delta) \sum_{(\eta,\xi^+)\in\Omega_M} f^M_t(\eta,\xi^+|d,\Delta)\omega_t(d,\Delta) \tag{3.4a}$$

$$+ \sum_{(d,h)\in\bar{S}} \pi(d,h) \sum_{(\eta,\,\xi^-)\in\Omega_M} f^M_t(\eta,\xi^-|d,h)V^*_{t'}(\boldsymbol{\pi'}) \tag{3.4b}$$

$$+ \sum_{(d,\Delta)\in\bar{\mathcal{S}}} \pi(d,\Delta) \sum_{(\eta,\xi^+)\in\Omega_M} f^M_t(\eta,\xi^+|d,\Delta)\mathcal{P}^M_t(d,\Delta)V^*_{t+1}(\boldsymbol{\pi'}) \tag{3.4c}$$

$$+ \sum_{h\neq\Delta,(d,h)\in\bar{\mathcal{S}}} \pi(d,h) \sum_{(\eta,\xi^+)\in\Omega_M} f^M_t(\eta,\xi^+|d,h)\psi_t(d,h). \tag{3.4d}$$

The optimal value function for the mammography action is collection of the immediate rewards (3.4a), rewards from the next decision epoch (3.4b)-(3.4c) and the terminal reward incurred when the patient is not healthy and the test result is positive (3.4d). The immediate reward consists of disutility of mammography and the quality-adjusted life years accrued when the test result is positive. The rewards from the next decision epoch are dependent on the mammography result. If the mammography result is negative, then the physician makes a decision about the supplemental screenings and the corresponding rewards are accrued. If the mammography result is positive, then the patient advances to the next decision epoch only if the test is false-positive.

In addition, we use the following optimal value functions to determine the supplemental screening decisions:

$$V_{t'}^{*}(\boldsymbol{\pi}) = \max_{a\in\{N\}\cup\mathcal{A}^{Sup}} \{V_{t'}^{a}(\boldsymbol{\pi})\}, \tag{3.5}$$

$$V_{t'}^{N}(\boldsymbol{\pi}) = r_t(\boldsymbol{\pi}) + \sum_{s\in\bar{\mathcal{S}}} \pi(s)\mathcal{P}_t^W(s)V_{t+1}^{*}(\boldsymbol{\pi'}), \tag{3.6}$$

$$\begin{aligned} V_{t'}^{a}(\boldsymbol{\pi}) = -\gamma(\boldsymbol{\pi}, a) &+ \sum_{(d,h)\in\bar{\mathcal{S}}} \pi(d,h) \sum_{(\eta,\ \xi^-)\in\Omega_a} f_t^a(\eta,\xi^-|d,h)\omega_t(d,h) \\ &+ \sum_{(d,h)\in\bar{\mathcal{S}}} \pi(d,h) \sum_{(\eta,\ \xi^-)\in\Omega_a} f_t^a(\eta,\xi^-|d,h)\mathcal{P}_t^a(d,h)V_{t+1}^{*}(\boldsymbol{\pi'}) \\ &+ \sum_{(d,\Delta)\in\bar{\mathcal{S}}} \pi(d,\Delta) \sum_{(\eta,\xi^+)\in\Omega_a} f_t^a(\eta,\xi^+|d,\Delta)\omega_t(d,\Delta) \\ &+ \sum_{(d,\Delta)\in\bar{\mathcal{S}}} \pi(d,\Delta) \sum_{(\eta,\xi^+)\in\Omega_a} f_t^a(\eta,\xi^+|d,\Delta)\mathcal{P}_t^a(d,h)V_{t+1}^{*}(\boldsymbol{\pi'}) \\ &+ \sum_{h\neq\Delta,(d,h)\in\bar{\mathcal{S}}} \pi(d,h) \sum_{(\eta,\xi^+)\in\Omega_a} f_t^a(\eta,\xi^+|d,h)\psi_t(d,h), \quad a\in\mathcal{A}^{Sup}. \end{aligned} \tag{3.7}$$

Note that the optimal value functions for a supplemental screening action $a \in \mathcal{A}^{Sup}$ is very similar to the optimal value function of the mammography action except that when the test result is negative, the patient accrues an immediate reward and advances to the next decision epoch. Finally, an optimal solution to our problem can be obtained by solving the following optimality equations:

$$V_t^{*}(\boldsymbol{\pi}) = \max\{V_t^W(\boldsymbol{\pi}),\ V_t^M(\boldsymbol{\pi})\}, \quad t<\mathcal{T},\ \boldsymbol{\pi}\in\Pi(\bar{\mathcal{S}}), \tag{3.8}$$

$$V_{\mathcal{T}}^{*}(\boldsymbol{\pi}) = \sum_{(d,h)\in\bar{S}} \pi(d,h)\omega_{\mathcal{T}}(d,h), \quad \boldsymbol{\pi}\in\Pi(\bar{\mathcal{S}}). \tag{3.9}$$

3.3 Solution Methodology

Current state-of-the-art in exact solution algorithms for POMDPs can only solve problems with very few core states. Therefore, several approximation methods are proposed to solve POMDP models more efficiently. Grid-based approaches are the most natural and widely-used approximations to solve POMDPs. In a recent study, Sandikci et al. [2013] use a fixed-resolution non-uniform grid method to solve their POMDP model for the liver allocation problem. Also, Zhang et al. [2012] approximate the solution to their POMDP model for prostate cancer screening problem by using a fixed-finite-grid method with a grid size of 1000. More information about the grid-based methods for POMDPs can be found in Lovejoy [1991a].

We approximate the optimal solution to the value functions (3.8)-(3.9) by discretizing belief space into a finite set of grid points. Let $\mathcal{G} = \{b_1, b_2, ..., b_{|\mathcal{G}|}\}$ be the grid set and $\mathcal{K} = \{1, ..., |\mathcal{G}|\}$ be the index set of $\mathcal{G}$. Assuming $\mathcal{G}$ includes all the extreme points of $\Pi(\bar{\mathcal{S}})$, any updated belief state $\boldsymbol{\pi'} \notin \mathcal{G}$ can be written as a convex combination of the points in $\mathcal{G}$ by using the convex interpolation

$$\boldsymbol{\pi'} = \sum_{k \in \mathcal{K}} \beta_k \boldsymbol{\ell}^k,$$

where $\boldsymbol{\ell}^k \in \mathcal{G}$, $\sum_{k=1}^{|\mathcal{G}|} \beta_k = 1,\ \beta_k \geq 0,$ for $k \in \mathcal{K}$ [Sandikci et al., 2013]. After constructing $\mathcal{G}$, we approximate the optimal value function $V_t^*(\cdot)$ by $\hat{\mathcal{V}}_t(\cdot)$, where $\hat{\mathcal{V}}_t(\boldsymbol{\pi'}) = \sum_{k \in \mathcal{K}} \beta_k \hat{\mathcal{V}}_t(\boldsymbol{\ell}^k),\ \forall \boldsymbol{\pi'} \notin \mathcal{G}$. Lovejoy [1991b] shows that this approximation provides an upper bound on the optimal solution as long as the corner points of the belief simplex $\Pi(\bar{\mathcal{S}})$ are included in $\mathcal{G}$.

3.3.1 Constructing the Grid Set

In any grid-based approach, the grid points are chosen such that all of the nonzero values in a grid point are positive integer multiples of $1/y$, where y is a positive integer representing the grid resolution. For instance, when the number of core states $|\bar{\mathcal{S}}| = 3$ and $y = 2$, we can form a

grid set consisting of following six grid points:

$$[1,0,0],\ [1/2,1/2,0],\ [1/2,0,1/2],\ [0,1,0],\ [0,1/2,1/2],\ [0,0,1].$$

A commonly used approach for generating a grid set is the fixed-resolution uniform grid method. For our problem, the number of grid points generated using the fixed-resolution uniform grid approach with resolution value y can be calculated as follows:

$$|\mathcal{G}_{\bar{\mathcal{S}}}| = \binom{|\bar{\mathcal{S}}| + y - 1}{y} = \frac{(|\bar{\mathcal{S}}| + y - 1)!}{y!(|\bar{\mathcal{S}}| - 1)!}. \tag{3.10}$$

Let $|\bar{\mathcal{S}}| = 10$ and resolution value $y = 100$. Then, the corresponding grid set contains approximately 4.3×10^{12} grid points. As the fixed-resolution uniform grid approach leads to large grid sets for increased resolution values, we use a variable-resolution uniform grid approach to form $\mathcal{G}_{\bar{\mathcal{S}}} \subset \Pi(\bar{\mathcal{S}})$. More specifically, we construct a grid set for partially observable health states ($\mathcal{G}_{\mathcal{C}} \subset \Pi(\mathcal{C})$), and another grid set for breast density states ($\mathcal{G}_{\mathcal{D}} \subset \Pi(\mathcal{D})$). Then, we construct the final grid set $\mathcal{G}_{\bar{\mathcal{S}}}$ by taking the outer product of the elements of each grid set as follows:

$$\mathcal{G}_{\bar{\mathcal{S}}} = \{v^i \otimes v^j,\ \forall v^i \in \mathcal{G}_{\mathcal{C}},\ \forall v^j \in \mathcal{G}_{\mathcal{D}}\}.$$

We use a different type of variable-resolution grid method to construct the grid set for the partially observable health states. We observe that the majority of the belief states included in $\Pi(\mathcal{C})$ are not representative of the breast cancer risk. For instance, no patients would have a health belief state corresponding to zero probability of being healthy unless they are diagnosed with breast cancer. Therefore, we focus on parts of $\Pi(\mathcal{C})$ that contain belief states that are potentially more representative of the patient's health status. For this purpose, we partition the $\Pi(\mathcal{C})$ into $\mathcal{I}$ regions and each region is described according to the first components of the belief states as

follows:

$$\kappa^i = \{\pi \in \Pi(\mathcal{C}) \mid \pi(0) \in [\varphi_{i-1}, \varphi_i)\},$$

where $\varphi = (\varphi_0, \varphi_1, ..., \varphi_{\mathcal{I}})$ is the vector of break points with $\varphi_0 = 1,$ and $\varphi_{\mathcal{I}} = 0$. We observe that the patients are more likely to have belief states for which the probability of being healthy is significantly higher than the probability of having cancer. Therefore, we use higher resolution values to sample from the regions that contain such belief states. That is, $y_1 \geq y_2 \geq \cdots \geq y_{\mathcal{I}}$, where y_i is the resolution value assigned to region κ^i. Let $\bar{\kappa}^i$ represent the union of the first i regions, i.e., $\bar{\kappa}^i = \cup_{j=1}^{i} \kappa^j$ and $\kappa^i = \bar{\kappa}^i \setminus \bar{\kappa}^{i-1}$. We obtain the number of grid points in $\bar{\kappa}^i$ using a resolution value of y_i as follows:

$$\phi(\bar{\kappa}^i, y) = \binom{|\mathcal{C}| + \lfloor (1 - \varphi_i)y \rfloor - 1}{|\mathcal{C}| - 1}.$$

Then, the total number of grid points for the variable-resolution grid approach is

$$|\mathcal{G}_{\mathcal{C}}| = \phi(\bar{\kappa}^1, y_1) + \sum_{i=2}^{\mathcal{I}} \Big(\phi(\bar{\kappa}^i, y_i) - \phi(\bar{\kappa}^{i-1}, y_i) \Big). \tag{3.11}$$

An example resolution assignment for the belief simplex $\Pi(\mathcal{C})$, when $|\mathcal{C}| = 3$ is given in Figure 3.2. In this example, we divide the belief space into four regions and use different resolution values to sample from each region. Let $y_1 = 1000,\ y_2 = 200,\ y_3 = 20,\ y_4 = 10$ be the resolution values that are used to sample grid points from region one through region four. Then, the variable-resolution uniform grid approach leads to a grid set with 1083 grid points. On the other hand, for the resolution value of $y = 1000$, the fixed-resolution uniform grid approach leads to a grid set with 501,501 grid points. That is, we can obtain the same accuracy at the focused region (i.e., region one) by using approximately 99.8% less points compared to the fixed-resolution uniform grid approach.

Next, we form the grid set $\mathcal{G}_{\mathcal{D}} \subset \Pi(\mathcal{D})$ for the density states. We use a fixed-resolution

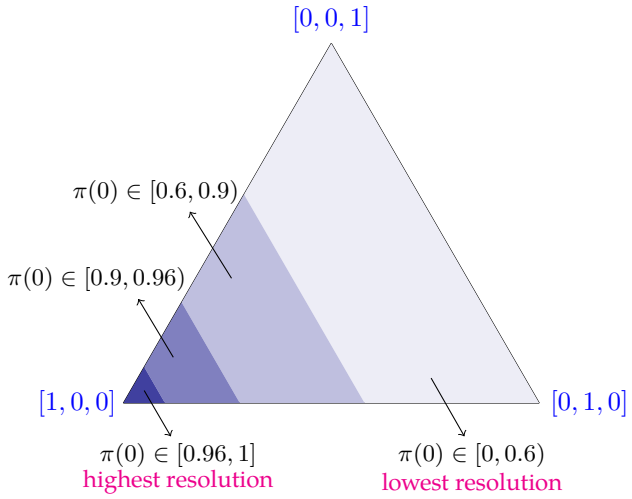

Figure 3.2: An example resolution setting for $\Pi(\mathcal{C})$

uniform grid approach to construct $\mathcal{G}_{\mathcal{D}}$. However, we eliminate the grid points that have nonzero values interleaved with zeros from the grid set. These grid points imply that the patient is assigned a nonzero probability of having breast density d_i and d_k, but zero probability of having breast density d_j, where $d_i < d_j < d_k$. Such grid points represent unlikely breast density probability distributions for a patient and it is not useful to include those points in the grid set. As a result, number of grid points in set $\mathcal{G}_{\mathcal{D}}$ reduces to [Sandikci et al., 2013]:

$$\sum_{i=0}^{min\{y-1,|\mathcal{D}|\}} \binom{y-1}{i} \cdot (|\mathcal{D}| - i).$$

3.3.2 Finding Interpolation Weights

There may be more than one set of interpolation weights, $\{\beta_k\}$, that spans the updated belief state, $\boldsymbol{\pi'}$, and choosing the combination that yields better approximate values (closer to $V^*(\pi')$) is important. We consider following approaches for obtaining interpolation weights:

1. *Iterative with full grid set* [IF]: In order to find the interpolation weights, Sandikci et al. [2013]

propose solving the following linear programming (LP) model:

$$\zeta(\mathcal{G}, \boldsymbol{\pi'}) = \min\Big\{ \sum_{k\in\mathcal{K}} \beta_k \hat{\mathcal{V}}_{t+1}(\boldsymbol{\ell}^k) \;\Big|\; \sum_{k\in\mathcal{K}} \beta_k \boldsymbol{\ell}^k = \boldsymbol{\pi'},\ \sum_{k\in\mathcal{K}} \beta_k = 1,\ \beta_k \geq 0\ \forall k \in \mathcal{K}\Big\}. \quad (3.12)$$

Observe that, to obtain the interpolation weights (β), this LP is solved at every decision epoch, for each action, for each observation and for each grid point, since they all may affect the value of updated belief state. Note that this approach iteratively calculates the β-values as the approximate QALY values for the grid points at decision epoch $t+1$ ($\hat{\mathcal{V}}_{t+1}(\boldsymbol{\ell}^k),\ \boldsymbol{\ell}^k \in \mathcal{G}$) are used in solving the LP.

2. *Iterative with disaggregated grid set* [ID]: The difficulty of computing $\zeta(\mathcal{G}, \boldsymbol{\pi'})$ increases in the size of the grid set. As an alternative, we use a subset of grid points, $\bar{\mathcal{G}} \subset \mathcal{G}$, to interpolate $\boldsymbol{\pi'}$, where $\bar{\mathcal{G}}$ is formed by using the corner points of the grid set for the density states $\mathcal{G}_{\mathcal{D}}$. That is,

$$\bar{\mathcal{G}} = \Big\{ v^i \otimes v^j,\ \forall v^i \in \mathcal{G}_{\mathcal{C}},\ \forall v^j \in \{e_1, ..., e_{|\mathcal{D}|}\}\Big\},$$

where $\boldsymbol{e}_k$ is a unit vector of size $|\mathcal{D}|$ with kth element equal to one and all other elements equal to zero. Note that, $\zeta(\bar{\mathcal{G}}, \boldsymbol{\pi'})$ gives an upper bound on the optimal value of $\zeta(\mathcal{G}, \boldsymbol{\pi'})$.

We can partition the set $\bar{\mathcal{G}}$ into smaller grid sets to calculate $\zeta(\bar{\mathcal{G}}, \boldsymbol{\pi'})$ more efficiently. Let $\bar{\mathcal{G}}_d$ represent the set of grid points that have non-zero entries for only breast density state d, i.e., $\bar{\mathcal{G}}_d = \{v \otimes e_d,\ \forall v \in \mathcal{G}_{\mathcal{C}}\}$, and $\mathcal{K}_d = \{1, 2, ..., |\bar{\mathcal{G}}_d|\}$ be the index set of this grid set. Then, $\bar{\mathcal{G}} = \cup_{d=1}^{\mathcal{D}} \bar{\mathcal{G}}_d$ and $\zeta(\bar{\mathcal{G}}, \boldsymbol{\pi'}) = \sum_{d=1}^{\mathcal{D}} \zeta_d(\bar{\mathcal{G}}_d, \boldsymbol{\pi'})$, where $\zeta_d(\bar{\mathcal{G}}_d, \boldsymbol{\pi'})$ can be described as follows:

$$\begin{aligned}\zeta_d(\bar{\mathcal{G}}_d, \boldsymbol{\pi'}) = \Big\{ \min \sum_{k\in\mathcal{K}_d} \beta_k \hat{\mathcal{V}}_{t+1}(\boldsymbol{\ell}^k) \;\Big|\; & \sum_{k\in\mathcal{K}_d} \beta_k \ell^k(d,h) = \pi'(d,h), \quad \forall h \in \mathcal{C}, \\ & \sum_{k\in\mathcal{K}_d} \beta_k = \sum_{h\in\mathcal{C}} \pi'(d,h), \\ & \beta_k \geq 0,\ \forall k \in \mathcal{K}_d\Big\}. \end{aligned} \quad (3.13)$$

As such, partitioning the grid set $\bar{\mathcal{G}}$ leads to a more efficient approach for obtaining the interpolation weights as we solve much smaller LPs.

3. *Distance-based* [DB]: Alternative to the iterative approaches, we can also choose the grid points used in the interpolation according to their distance to the updated belief state $\boldsymbol{\pi'}$. More specifically, interpolation weights can be obtained using the following LP:

$$\Upsilon(\bar{\mathcal{G}}, \boldsymbol{\pi'}) = \min \Big\{ \sum_{k \in \mathcal{K}} \sum_{s \in \bar{\mathcal{S}}} \beta_k |\pi'(s) - \ell^k(s)| \;\Big|\; \sum_{k \in \mathcal{K}} \beta_k \boldsymbol{\ell}^k = \boldsymbol{\pi'},\ \sum_{k \in \mathcal{K}} \beta_k = 1,\ \beta_k \geq 0 \ \forall k \in \mathcal{K} \Big\}, \tag{3.14}$$

where $\bar{\mathcal{G}} \subseteq \mathcal{G}$. While we use Manhattan distance as the distance measure in above LP, other metrics can also be employed for $\Upsilon(\bar{\mathcal{G}}, \boldsymbol{\pi'})$. Furthermore, $\bar{\mathcal{G}}$ can be refined to reduce the size of the LP. That is, we can construct $\bar{\mathcal{G}}$ using n closest points to $\boldsymbol{\pi'}$ and the corner points of $\Pi(\bar{\mathcal{S}})$, so that LP is always feasible. Also note that distance-based approach can be extended to [ID] as well (i.e., when $\bar{\mathcal{G}} = \cup_{d=1}^{\mathcal{D}} \bar{\mathcal{G}}_d$).

While the LPs can be solved in parallel for a given decision epoch for all three approaches, [DB] is more advantageous in terms of parallelization as we can calculate the interpolation weights for each period independently. Furthermore, iterative approaches can take prohibitively long times as the size of $\mathcal{G}$ increases. For instance, when we use a grid set with approximately 4400 grid points for our problem (using the parameter setting given in Section 3.4) [IF], [ID], and [DB] takes 18.0, 6.0, and 1.5 hours to calculate the β-values, respectively. In addition, difference in the resulting approximate function values for the grid points is less than %0.01.

3.4 Parameter Estimation

In our numerical experiments, we define each decision epoch to be one year and consider patients between ages 40 and 100; hence $\mathcal{T} = 60$. Health status of a patient is generally categorized into three states in terms of breast cancer, namely, cancer-free, in situ cancer, and invasive cancer

[Fryback et al., 2006]. Therefore, we take $\mathcal{C} = \{\Delta, h_1, h_2\}$, where h_1 corresponds to in situ cancer, and h_2 corresponds invasive cancer. Note that in situ cancer is an early stage cancer and can become invasive over time whereas invasive cancer corresponds to a late stage cancer. We use the categorization suggested by the Breast Imaging Reporting and Data System (BI-RADS) for breast density. In particular, we take $\mathcal{D} = \{d_0, d_1, d_2, d_3\}$, where d_0 corresponds to *fatty*, d_1 corresponds to *scattered fibrodengular*, d_2 corresponds to *heterogeneously dense* and d_3 corresponds to *extremely dense* breast density categories.

We only consider MRI and ultrasound as supplemental screening modalities since other modalities such as tomosynthesis and molecular breast imaging are not widely used in breast cancer screening. MRI uses magnets and radio waves to produce detailed 3-dimensional images of the breast tissue. The patient has a contrast solution injected before the MRI and she needs to remain still during the test for an extended amount of time, which causes a lot of stress and discomfort. On the other hand, ultrasound is a significantly less invasive procedure and the test is generally performed using a handheld device that uses sound waves to make images of the breast. Unlike mammography, neither MRI nor ultrasound exposes the patient to any radiation.

We use several previously validated data sources to estimate the parameters of our model. The validity of our model and resulting policy insights are based on these data sources and our comparisons with the relevant studies in the literature. Our primary sources of data for the parameter estimation are the University of Wisconsin Breast Cancer Simulation (UWBCS) and the Surveillance, Epidemiology, and End Results (SEER) Program. More specifically, we estimate the transition probabilities using UWBCS and post-cancer lifetime from SEER. We present the details of parameter estimation regarding screening exam performance in Section 3.4.1, the disutility values associated with the screening exams in Section 3.4.2, and breast cancer risk estimation in Section 3.4.3.

3.4.1 Performance of the screening exams

In a given decision epoch, if the patient does not receive any type of screening exam, she can make a self-detection either via clinical breast examination or breast self exam. We use Barton et al. [1999]'s and Baxter et al. [2001]'s studies to estimate sensitivity and specificity of the clinical breast exams and breast self exam, respectively. We obtain the proportion of clinical breast exam and breast self exam from Elmore et al. [2005] and Messina et al. [2004] to calculate the sensitivity and specificity of the self-detection by a patient. Table 3.1 summarizes the sensitivity and specificity values for breast self-exam (BSE), clinical breast exam (CBE), and self-detection.

BSE, CBE, and self-detection			
Test	Sens. (%)	Spec. (%)	Data source
BSE	26	90	Baxter et al. [2001]
CBE	54	94	Barton et al. [1999]
Self-detection	44	92	-

Table 3.1: Performance values for the Wait action.

Table 3.2 shows the sensitivity and specificity values for different screening modalities. Mammography sensitivity and specificity values are estimated from Sprague et al. [2014]'s study. Note that digital mammography performs best for breast density d_1 and performs worst for breast density d_0, whereas film mammography performance decreases with increased breast density. We also obtain the supplemental ultrasound sensitivity and specificity values from Sprague et al. [2014]. In our experiments, we use the ultrasound sensitivity and specificity as 55% and 94%, respectively. Several studies report that MRI sensitivity varies between 88-95%, and MRI specificity varies between 60-75% [Peters et al., 2008, Knuttel et al., 2014]. As such, we use the values for MRI sensitivity and specificity as 90% and 72%, respectively. Both MRI and ultrasound performance are not affected by age and breast density of the patient.

Density	Age group	Film		Digital	
		Sensitivity (%)	**Specificity (%)**	**Sensitivity (%)**	**Specificity (%)**
d_0	40-49	81.6	95.5	71.3	94.2
	50-74	86.3	96.2	77.9	95.2
	75-100	91.5	96.2	83.1	95.2
d_1	40-49	79.9	91.6	83.6	89.5
	50-74	85.3	92.9	88	91.2
	75-100	90.5	92.9	93.2	91.2
d_2	40-49	75.0	89.2	75.9	86.4
	50-74	80.7	90.2	81.6	88.5
	75-100	85.9	90.2	86.8	88.5
d_3	40-49	59.7	90.4	76.0	89.3
	50-74	67.6	92.0	81.7	91.1
	75-100	72.8	92.0	86.9	91.1

Table 3.2: Mammography screening performance by density and age group.

3.4.2 Disutility values for the screening exams

Table 3.3 shows the disutility values associated with screening exams and positive test results. Several studies report that patients are more cautious towards MRI when compared to mammography or ultrasound, as MRI is a longer and more painful exam [Berg et al., 2009]. Therefore, disutility of MRI exam is higher compared to mammography and ultrasound. In our analyses, we take the disutility of undergoing MRI as two days and we explore sensitivity with respect to this parameter in Section 3.5.3.

3.4.3 Personal breast cancer risk

Several risk estimation tools are developed to predict the breast cancer risk of a patient. In our analysis, we use the BCSC risk calculator to estimate the breast cancer risk for individual patients (https://tools.bcsc-scc.org/BC5yearRisk/calculator.htm). Table 3.4 summarizes the

Intervention	Disutility values	Data source
Self-detection	0 days	-
Mammography	0.5 days	Mandelblatt et al. [1992]
Ultrasound	0.5 days	-
MRI	2 days	-
False Positive	2 weeks	Chhatwal et al. [2010]
True Positive	2 weeks	Velanovich [1995]

Table 3.3: Interventions and associated QALY decrements.

invasive cancer risk estimates for 40-year old patients according to their risk categories and breast densities.

	d_0	d_1	d_2	d_3
average-risk	0.08%	0.17%	0.29%	0.35%
high-risk	0.17%	0.36%	0.61%	0.74%
very high-risk	0.55%	1.19%	1.99%	2.44%

Table 3.4: Invasive cancer risk at age 40 for patients with different breast densities and risk levels.

In our analysis, we categorize the women into three risk groups: average-risk, high-risk, and very high-risk. Our risk estimates suggest that, depending on the breast density, a high-risk woman is approximately two times more likely to get cancer and a very high-risk woman is approximately six to seven times more likely to get cancer compared to an average risk woman. It is important to note that BCSC risk calculator does not include some risk factors such as BRCA1 and BRCA2 gene mutations and having previous chest radiation. However, such risk factors can be easily incorporated into our model. For instance, women with BRCA1 gene mutations are approximately seven to eight times more likely to get breast cancer [Chen and Parmigiani, 2007], and by adjusting our initial breast cancer risk estimates and transition probability matrices accordingly, we can make our model suitable for women with BRCA1 gene mutations.

3.5 Numerical Results

We solve the approximation of our POMDP model using the backward induction algorithm. In our base case grid settings, we partition the health belief space, $\Pi(\mathcal{C})$, into four regions according to the first components of the health belief states. We use the vector of break points $\varphi = (1, 0.96, 0.9, 0.6, 0)$ to define the regions (e.g., first region is defined by $\kappa^1 = \{\pi \in \Pi(\mathcal{C}) | \pi(0) \in [1, 0.96)\}$). Also, we use the resolution values 1000, 200, 20 and 10 to sample grid points from $\kappa^1,\ \kappa^2,\ \kappa^3,$ and κ^4, respectively. This resolution setting leads to a health grid set ($\mathcal{G}_\mathcal{C}$) with 1083 grid points. Considering that we only need the corner points of the breast density belief space, our final grid set contains 4332 grid points. We observe that this grid setting provides tight upper bounds for the QALY values (see Appendix A.3 for a detailed performance analysis).

3.5.1 Optimal Mammography Screening Policies

Figure 3.3 shows the in situ and invasive cancer risk combinations for which it is optimal to recommend each screening modality when film mammography is used as the initial screening test. We observe that ultrasound is more preferable over MRI when breast cancer risk is relatively low. Furthermore, risk combinations to recommend each action changes based on the patient's breast density. In particular, patients are recommended more supplemental screenings as breast density increases.

We do not observe the same monotone increase in supplemental screening recommendations with increased breast density when digital mammography is used as the initial screening test (see Figure 3.4). This is mainly due to nonuniform performance of digital mammography across the breast densities. For instance, considering that the digital mammography performs best for the patients with breast density d_1, the area of risk combinations for which supplemental screenings are recommended is smallest for d_1. As digital mammography is more widely used in the clinics [Stout et al., 2014], we use digital mammography as the initial screening test in the rest of the numerical experiments.

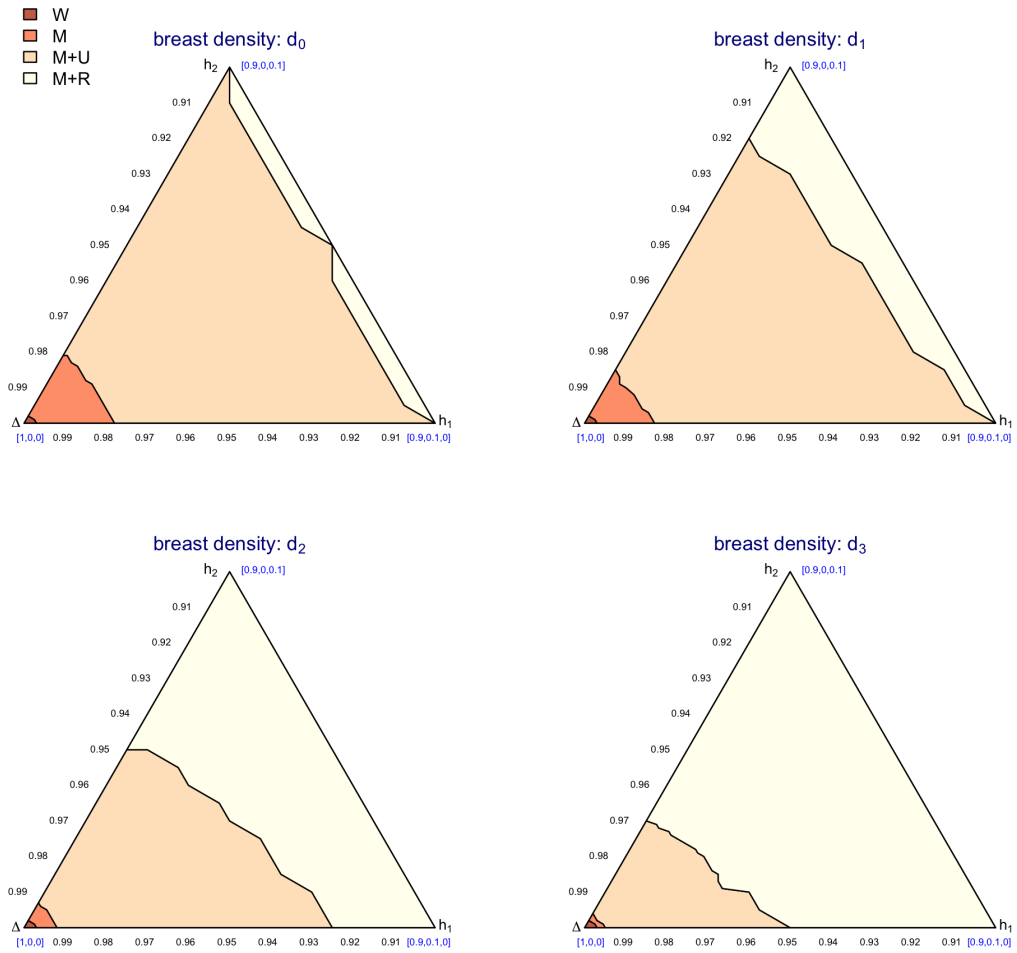

Figure 3.3: Film mammography risk thresholds for a 40-year old average-risk patient

An important aspect of personalized screening recommendations is that a patient's screening history affects her future screening recommendations. We consider two possible screening scenarios for a patient and report the number of mammography and supplemental screenings recommended for these scenarios in Table 3.5. "$Sce1$" corresponds to the scenario where the patient always makes negative observations as a result of any action and "$Sce2$" corresponds to the case where the patient makes positive observations due to wait action every 10 years. If the patient is in the average-risk category and her breast density is d_0 or d_1, she is not recom-

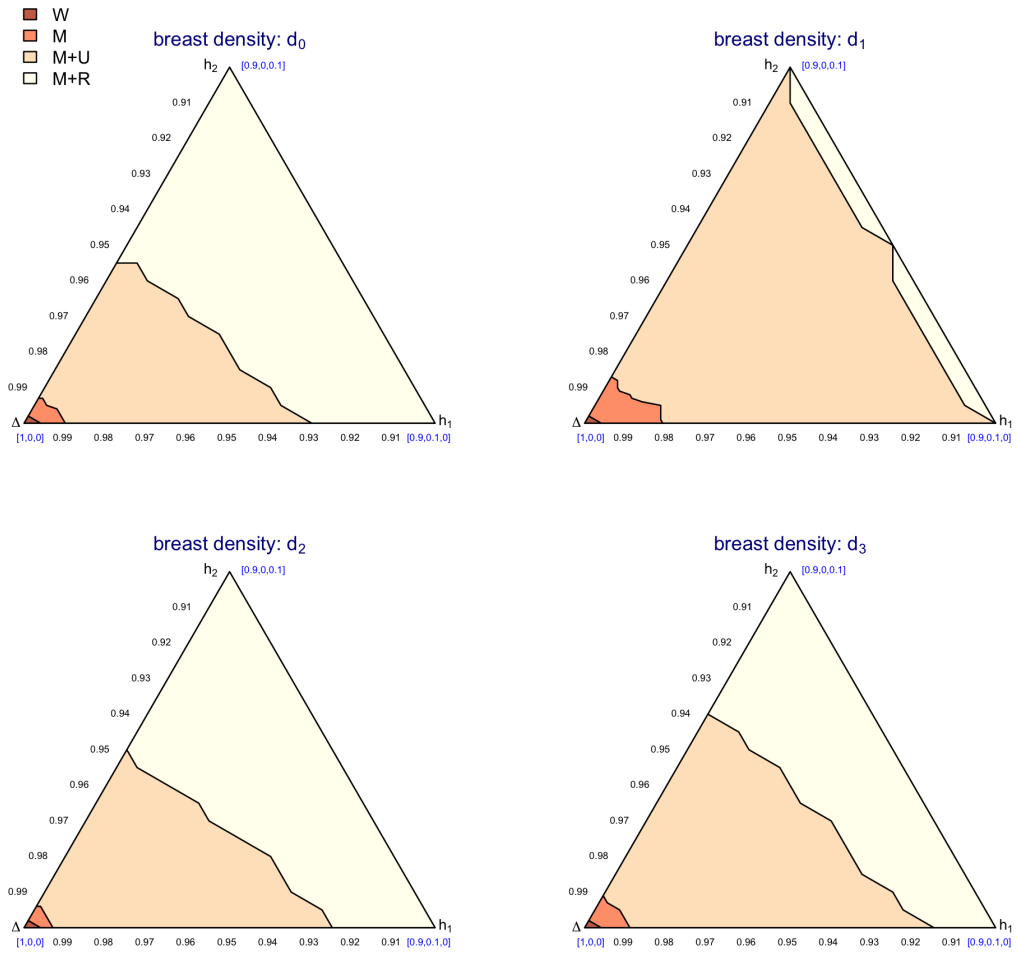

Figure 3.4: Digital mammography risk thresholds for a 40-year old average-risk patient

mended supplemental screening tests for both scenarios. That is, increased breast cancer risk due to positive observations would only lead to increased number of mammography screening recommendations but not supplemental screening recommendations. On the other hand, if the patient is in the high-risk category or has dense breasts (d_2 or d_3), she might be recommended a limited number of ultrasound screenings.

		average-risk			high-risk			very high-risk		
		M	R	U	M	R	U	M	R	U
d_0	Sce1	9	0	0	27	0	1	39	0	1
	Sce2	16	0	1	29	0	2	39	0	2
d_1	Sce1	15	0	0	27	0	0	40	0	1
	Sce2	18	0	0	30	0	1	40	0	1
d_2	Sce1	16	0	0	38	0	1	40	0	6
	Sce2	19	0	3	39	0	2	40	0	6
d_3	Sce1	24	0	0	40	0	0	40	0	18
	Sce2	26	0	2	40	0	0	40	0	18

Table 3.5: Number of screenings recommended for patients from different risk groups under different scenarios.

3.5.2 Value of Modeling Breast Density and Supplemental Screenings

We next compare our POMDP model to population based screening guidelines that have been in effect in many countries. Note that these guidelines do not make any supplemental screening recommendations. Table 3.6 shows the results of this experiment for very high-risk patients (see Appendix A.4 for the experiments regarding average-risk and high-risk patients). Each screening guideline is given by the frequency of screening as well as the screening start and end ages. For example, "tri_50_70" represents triennial screening between ages 50 and 70. Also, "noScr" represents the case where the patient does not receive any screenings, "pomdp" represents our model and "pomdp_nss" represents a modified version of our model, where patients are not allowed to receive supplemental screenings. We compare these approaches according to expected number of screenings recommended (avg. Scr), expected number of false-positive screenings (avg. FP) and QALY values for the patient (QALY) where "QALY gain" column shows the QALY improvements over the "noScr" scenario in months.

Dynamic screening policies obtained by our model (pomdp and pomdp_nss) outperform population-based screening guidelines in terms of QALY gains. Also, pomdp models provide patient-specific screening recommendations, therefore prevents unnecessary screenings and

	avg. Scr	avg. FP	QALY	QALY gain (months)
noScr	0.00	0.00	36.712	-
tri_50_70	9.16	0.41	38.212	18.00
bi_50_70	12.77	0.64	38.418	20.47
bi_40_74	19.94	1.08	39.096	28.60
an_40_74	34.06	2.09	39.226	30.18
pomdp_nss	37.68	2.35	39.428	32.59
pomdp	37.82	2.37	39.437	32.70

(a) Breast density: d_0

	avg. Scr	avg. FP	QALY	QALY gain (months)
noScr	0.00	0.00	34.684	-
tri_50_70	11.25	0.56	37.330	31.75
bi_50_70	14.49	0.86	37.592	34.89
bi_40_74	21.75	1.49	38.477	45.52
an_40_74	34.92	2.90	38.652	47.62
pomdp_nss	37.90	3.18	38.865	50.18
pomdp	38.70	3.24	38.868	50.22

(b) Breast density: d_1

	avg. Scr	avg. FP	QALY	QALY gain (months)
noScr	0.00	0.00	33.086	-
tri_50_70	12.60	0.58	36.539	41.43
bi_50_70	15.96	0.89	36.859	45.28
bi_40_74	22.99	1.65	37.894	57.70
an_40_74	35.46	3.20	38.164	60.94
pomdp_nss	37.92	3.48	38.369	63.40
pomdp	39.98	3.66	38.377	63.50

(c) Breast density: d_2

	avg. Scr	avg. FP	QALY	QALY gain (months)
noScr	0.00	0.00	31.955	-
tri_50_70	13.69	0.53	36.003	48.58
bi_50_70	16.64	0.80	36.354	52.79
bi_40_74	24.08	1.44	37.526	66.86
an_40_74	35.94	2.84	37.866	70.93
pomdp_nss	37.94	3.06	38.058	73.24
pomdp	45.23	3.66	38.072	73.40

(d) Breast density: d_3

Table 3.6: Comparison of screening policies for very high-risk patients

false-positive test results. We observe that supplemental screenings leads to marginal QALY gains for every breast density category, which indicates that they are ineffective in increasing QALY of the patients. Also note that breast density of the patient significantly affects the screening recommendations. For example, our POMDP model (pomdp) leads to 37.82 screening recommendations (on average) with 32.70 months QALY gain for the patient with d_0 breast density, and 45.23 screening recommendations (on average) and 73.40 months QALY gains for the patient with d_3 breast density.

Table 3.7 shows the expected number of screenings for each breast density and risk-group categories. As the supplemental screenings are only allowed after negative mammography result, we also present the percentage of time each supplemental screening method is recommended. For example, a very high-risk patient with breast density d_3 is recommended to undergo ultrasound

33.8% of the time she receive a negative mammogram. These results show that the patients are infrequently recommended to receive supplemental screenings. Also, we observe that increased breast cancer risk mainly leads to more mammography screening recommendations rather than increasing the amount of supplemental screenings recommended.

	average-risk			high-risk			very high-risk		
	E[M]	E[R]	E[U]	E[M]	E[R]	E[U]	E[M]	E[R]	E[U]
d_0	13.20	0(0.0%)	0.01(0.0%)	26.38	0(0.0%)	1.11(4.9%)	36.83	0(0.0%)	0.99(3.5%)
d_1	15.10	0(0.0%)	0.02(0.2%)	25.47	0(0.0%)	0.20(1.0%)	37.83	0(0.0%)	0.87(3.4%)
d_2	15.69	0(0.0%)	0.72(6.2%)	31.44	0(0.0%)	1.05(4.5%)	37.94	0(0.0%)	2.03(8.9%)
d_3	19.61	0(0.0%)	0.73(5.0%)	35.66	0(0.0%)	0.06(0.2%)	38.00	0(0.0%)	7.23(33.8%)

Table 3.7: Number of screenings recommended.

3.5.3 Sensitivity Analysis

We conduct a one-way sensitivity analysis on sensitivity and specificity of the screening modalities to demonstrate the impact of test performance on the number of screening recommendations. In this experiment, we only consider a very-high risk patient and assume that the patient does not receive any positive observations throughout her lifetime. Table 3.8 shows the range of screening counts by screening modality based on the patient's breast density. For example, in Table 3.8(a), number of mammography screenings recommended to the patient with d_0 breast density changes between 38 and 37 as the sensitivity of the mammography changes. Sensitivity analysis for the specificity values indicate that the number of screenings recommended to a patient is significantly affected by the variations in the specificity of the tests. For example, a 5% decrease in the specificity of mammography would lead to recommending 20 ultrasounds to a patient with d_3 breast density, whereas increasing the mammography specificity by 5% leads to 13 ultrasound recommendations for the same patient.

One-way sensitivity analysis can also be used to investigate the concerns about the performance of the supplemental tests. For example, a main drawback of MRI is its low specificity.

Table 3.8(b) shows that even a 20% increase in the specificity value of MRI would not increase the number of MRI recommendations. Similarly, ultrasound is known to have low sensitivity. Table 3.8(c) shows that there is a significant increase on the number of ultrasound recommendations to the patients with dense breasts as the ultrasound sensitivity improves. We also conduct a one-way sensitivity analysis for the disutility of MRI by varying the disutility value between 0.5 days and two days (base case value is two days). However, we do not observe any change in the results with decreased disutility values for MRI, which strengthen our insights that the main factor for fewer MRI recommendations is its low specificity.

3.6 Concluding remarks

We investigate the breast cancer screening problem considering the breast density of the patient and availability of the supplemental screenings. We formulate this problem as a POMDP and propose a grid-based approximation to solve the POMDP model. We use breast cancer epidemiology data in our numerical experiments and observe that breast density information could significantly affect screening recommendations. In addition, our numerical analysis indicates that for the patients with higher breast cancer risk, more frequent mammography screenings are preferable over supplemental screenings. As such, our study makes a potential contribution to the recent debate about breast density legislation that recommends patients with dense breasts to discuss supplemental screening tests with their providers.

Replacement of film mammography by digital mammography in the United States has significantly reduced the difference in screening mammography sensitivity according to breast density. In particular, sensitivity of digital mammography for women with dense breasts is around 76% to 82% depending on age. Therefore, benefits of supplemental screenings diminish as digital mammography already performs well for women with dense breasts. In this regard, our findings mostly overlaps with two recent studies by Sprague et al. [2014] and Melnikow et al. [2016] that analyze the impact of supplemental screenings. On the other hand, although we

	$sens(M) \in [-0.05, +0.05]$				$spec(M) \in [-0.05, +0.05]$			
	M	R	U	QALY change	M	R	U	QALY change
d_0	(38,37)	(1,0)	(0,1)	0.564	(38,39)	(0,0)	(1,1)	1.478
d_1	(40,40)	(0,0)	(1,1)	0.692	(40,40)	(0,0)	(1,1)	1.338
d_2	(40,40)	(0,0)	(26,1)	0.796	(40,40)	(0,0)	(11,4)	1.193
d_3	(40,40)	(0,0)	(33,2)	0.768	(40,40)	(0,0)	(20,13)	1.078

(a) One-way sensitivity analysis for mammography

	$sens(R) \in [-0.05, +0.05]$				$spec(R) \in [+0.05, +0.20]$			
	M	R	U	QALY change	M	R	U	QALY change
d_0	(39,39)	(0,0)	(1,1)	0.000	(39,39)	(1,1)	(0,0)	0.042
d_1	(40,40)	(0,0)	(1,1)	0.000	(40,40)	(0,0)	(1,1)	0.000
d_2	(40,40)	(0,0)	(6,6)	0.000	(40,40)	(0,1)	(6,5)	0.023
d_3	(40,40)	(0,0)	(18,18)	0.000	(40,40)	(0,1)	(18,16)	0.017

(b) One-way sensitivity analysis for MRI

	$sens(U) \in [+0.10, +0.30]$				$spec(U) \in [-0.05, +0.05]$			
	M	R	U	QALY change	M	R	U	QALY change
d_0	(38,39)	(0,0)	(1,1)	0.060	(40,39)	(1,0)	(0,1)	0.046
d_1	(40,40)	(0,0)	(1,1)	0.031	(40,40)	(0,0)	(1,1)	0.040
d_2	(40,40)	(0,0)	(17,28)	0.182	(40,40)	(0,0)	(1,38)	0.251
d_3	(40,40)	(0,0)	(24,34)	0.282	(40,40)	(0,0)	(2,37)	0.421

(c) One-way sensitivity analysis for ultrasound

Table 3.8: Sensitivity analysis considering an very high-risk patient

consider very high-risk women in our study, some women might have more than one strong risk factors such as BRCA1 gene mutations and multiple first-degree relatives who have had breast cancer. We recognize that such patients should not follow the general guidelines and should be screened more aggressively.

4 ANALYSIS OF MAMMOGRAPHY SCREENING POLICIES UNDER RESOURCE CONSTRAINTS

4.1 Problem Definition

In this chapter, we provide a framework for analyzing breast cancer screening policies both at the individual and population level. As a part of this framework, we maximize the total expected QALYs of the patients under resource constraints. Specifically, we develop a finite-horizon discrete-time constrained POMDP model to optimize the allocation of limited mammography screening resources to maximize the benefits of mammography screening. We reformulate the constrained POMDP model as a mixed-integer program and further extend this model to devise a resource allocation scheme between patients from different risk groups.

We impose a limit on the number of mammography screenings instead of the actual cost of screening, since the number of mammography screenings captures various resource requirements including the cost of screening, number of machines to take mammography exams, and the number of trained personnel to interpret mammograms. Note that the screening limits can be determined based on the available resources for a given country or a screening program. For instance, in Chile, the public health system provides only two mammography screenings throughout the lifetime of a woman [OECD, 2013].

Our work makes several important contributions to both theory and practice. First, as far as we are aware, ours is the first study to consider the resource constraints and imperfect test results simultaneously while optimizing screening decisions in a cancer screening problem. Second, we use a novel approach to obtain clinically intuitive policies. More specifically, the use of a mixed-integer linear program allows us to impose constraints to ensure that the resulting optimal policies are of the control-limit type, which are easy to implement and more appealing to policy makers. Third, using breast cancer epidemiology data, we test our approach and find important policy implications that would help policy makers in a resource-limited setting. For

example, we find that it is not cost-effective to allocate more than eight screenings to average-risk patients throughout their lives. Finally, the explicit modeling of the constraints allows us to overcome limitations of previous studies in health policy making where the subjective disutility values associated with the interventions may significantly affect the recommended policies. We numerically demonstrate that our proposed screening policies are robust to the disutility values associated with mammography screening.

The remainder of this chapter is organized as follows. Section 4.2 describes the model components and presents our constrained POMDP model. We introduce our approximation scheme to solve the constrained POMDP model in Section 4.3 and we formulate a mixed-integer linear programming model for resource allocation problem in Section 4.4. In Section 4.5, we describe our data and provide the details of parameter estimation. In Section 4.6, we report our numerical results. Finally, we summarize our findings and conclude in Section 4.7.

4.2 The POMDP Model

We develop a finite-horizon discrete-time constrained POMDP model to optimize the allocation of limited mammography screening resources to maximize the benefits of mammography screening. We aim to find a screening policy that maximizes a patient's QALYs given that there is a limit on the available number of mammography screenings. This screening policy determines the set of actions to be taken at each decision epoch. The decisions are made at every 6 months starting from age 40 (the earliest recommended starting age for screening) until 90. Therefore, the set of *decision epochs* is $\mathcal{T} = \{0, 1, 2, \ldots, N-1\}$, where $N = 100$.

The *state space* of the POMDP model consists of the health status of the patients. We use $s_t \in \mathcal{S}$ to represent the core states (i.e., underlying true health status of a patient at time t), where $\mathcal{S} = \{healthy,\ in\,situ,\ invasive,\ post\text{-}in\,situ,\ post\text{-}invasive,\ death\}$. Note that in situ state corresponds to an early stage cancer, whereas invasive state corresponds to a more advanced stage cancer. Furthermore, as mammography screening is not perfectly accurate, a patient's actual disease

state is only partially observed. Therefore, we partition the core states as partially observable and fully observable, and denote the partially observable states by $\bar{\mathcal{S}} = \{healthy,\ in\ situ,\ invasive\}$. For the brevity of notation, we match each state s with an integer, $i \in \{0, 1, ..., 5\}$, to represent the states in analytical expressions.

Our model optimizes the mammography screening decisions based on patients' personal risk of breast cancer. The personal risk of breast cancer corresponds to the decision maker's assessment of the patient's disease state and is summarized by a *belief state*, b, that is defined over the partially observable states. Because the core state of the patient is not always fully observable (i.e., when $s_t \in \bar{\mathcal{S}}$), belief state b represents a probability distribution over $\bar{\mathcal{S}}$ and the *belief space* consisting of all such probability distributions is defined as

$$\mathcal{B}(\bar{\mathcal{S}}) = \left\{ b \in [0,1]^{|\bar{\mathcal{S}}|} : \sum_{s \in \bar{\mathcal{S}}} b(s) = 1 \right\}, \tag{4.1}$$

where $b(s)$ corresponds to the probability of being in core state $s \in \bar{\mathcal{S}}$.

At each decision epoch, depending on the patient's health state, the decision maker can take one of the two types of *actions*: recommend that the patient undergo mammography screening (M) or do nothing (DN) and the patient waits until the next decision epoch. Hence, the action at decision epoch t is denoted by $a_t \in \mathcal{A} = \{DN, M\}$.

4.2.1 Observation Probabilities

As a result of the action taken, an *observation* is made about the patient's health state. If the action taken is DN, there can be a self-detection $(DN+)$, which is the case when the patient and/or her physician feels a lump in the breast, or no self-detection $(DN-)$. Similarly, if the action taken is M, the outcome of the mammography is either positive $(M+)$ or negative $(M-)$. Hence, we define the *observation space* as $O = O_{DN} \cup O_M$, where $O_{DN} = \{DN-, DN+\}$ and $O_M = \{M-, M+\}$.

We use $z_t(o_t|s_t, a_t)$ to represent the probability of observing o_t when the patient's health state is s_t and the action taken is a_t. Observation probabilities are obtained from the sensitivity and specificity of the associated action. For instance, mammography screening can detect in situ ($s = 1$) or invasive ($s = 2$) cancer with probabilities equal to its sensitivity, denoted as τ^s_M. Similarly, the patient can make a self-detection with probability τ^s_{DN}, if she is in state s. On the other hand, if the patient does not have cancer, each action may lead to a false-positive diagnosis with probability $1 - w_{a_t}$, where w_{a_t} is the specificity of the action taken. Table A.2 summarizes observation probabilities for different core state-observation pairs. Note that while sensitivity and specificity of each action is dependent on the patient's age, we drop the time index for notational clarity.

$s \backslash o$	$DN-$	$DN+$	$M-$	$M+$
$s = 0$	w_{DN}	$1 - w_{DN}$	w_M	$1 - w_M$
$s \in \{1, 2\}$	$1 - \tau^s_{DN}$	τ^s_{DN}	$1 - \tau^s_M$	τ^s_M

Table 4.1: Observation probabilities $\mathbf{z_t}(\mathbf{o_t}|\mathbf{s_t}, \mathbf{a_t})$

4.2.2 Transition Probabilities

The probability $p_t(s_{t+1}|s_t, a_t)$ denotes the *health state transition probability* from state s_t to s_{t+1} given that the patient takes action a_t at decision epoch t. That is,

$$p_t(s_{t+1}|s_t, a_t) = \sum_{o_t \in O_{a_t}} p_t(s_{t+1}|s_t, a_t, o_t) z_t(o_t|s_t, a_t),$$

where $p_t(s_{t+1}|s_t, a_t, o_t)$ corresponds to the transition probability when o_t is observed as a result of action a_t. Note that the probability of a change in health state in the next decision epoch is the same for women who have a negative or false-positive mammogram, and those who do not receive a mammogram. The relations between transition probabilities for different actions and

observations are as follows:

$$p_t(s_{t+1}|s_t, DN, DN-) = p_t(s_{t+1}|s_t, DN, DN+) = p_t(s_{t+1}|s_t, M, M-),$$

$$p_t(s_{t+1}|0, M, M-) = p_t(s_{t+1}|0, M, M+).$$

4.2.3 Rewards

The *immediate rewards* for our model correspond to the QALYs accrued in a decision epoch. QALY is a measure of expected difference between the lifetime and the disutility of the associated action and is commonly used in designing health policies [Gray et al., 2010]. More specifically, if the patient is in state s_t, takes action a_t and makes observation o_t, she accrues an immediate reward $r_t(s_t, a_t, o_t)$. We use half-cycle correction method to calculate the immediate reward associated with the DN action [Gray et al., 2010]. That is, we assume half decision interval length for the expected life in the case of death and full decision interval length when the patient survives until the next decision epoch. Then, $r_t(s_t, DN, \cdot)$ is calculated as the weighted average of expected life by weighing life years until the subsequent decision epoch according to the probability of death during that time frame, that is,

$$r_t(s_t, DN, \cdot) = 0.5 \times P(\text{alive in period } t|\text{current state is } s_t) + 0.25 \times P(\text{dies in period } t|\text{current state is } s_t).$$

We subtract the disutility of undergoing a mammography screening from $r_t(s_t, DN, \cdot)$ to calculate the immediate reward for the mammography action. That is,

$$r_t(s_t, M, o_t) = r_t(s_t, DN, \cdot) - \gamma^M(s_t, o_t), \quad \forall s_t \in \bar{\mathcal{S}},$$

where $\gamma^M(s_t, o_t)$ be the disutility associated with mammography screening when the patient's health state is s_t and the patient observes $o_t \in O_M$. Note that $\gamma^M(s_t, M-),\ s_t \in \bar{\mathcal{S}}$, corresponds

to disutility of negative mammogram, $\gamma^M(0, M+)$ corresponds to the disutility of a false-positive (FP) mammogram and $\gamma^M(s_t, M+),\ s_t \in \{1, 2\}$, corresponds to the disutility of a true positive (TP) mammogram. Furthermore, we can calculate the immediate rewards for each action and state as follows:

$$r_t(s_t, a_t) = \sum_{o_t \in O_{a_t}} r_t(s_t, a_t, o_t) z_t(o_t | s_t, a_t).$$

If the patient is in state $s_t \in \{1, 2\}$ and receives a positive mammography result, she starts treatment, quits the decision process, and accrues a lump-sum reward $R_t(s_t)$. Also, at the final decision epoch, no action is taken and the patient accrues a terminal reward $R_N(s)$, if she is in state s. Finally, we can calculate the expected immediate reward the patient accrues for taking action a_t in belief state b_t as

$$r_t(b_t, a_t) = \sum_{s_t \in \bar{\mathcal{S}}} b_t(s_t) r_t(s_t, a_t), \quad \text{and} \quad R_t(b_t) = \sum_{s_t \in \bar{\mathcal{S}}} b_t(s_t) R_t(s_t).$$

4.2.4 Optimality Equations

The optimal policy for the unconstrained breast cancer screening problem yields the set of actions that maximize a patient's total QALYs through the planning horizon. We obtain the optimal policy by dynamically solving the corresponding optimality equations. We formulate these optimality equations using the belief state, which preserves the Markovian property and also is a sufficient statistic for the entire history of the process [Smallwood and Sondik, 1973]. Let $Q_t^*(b_t)$ be the maximum total expected QALYs that the patient attains when she is in belief state

b_t at time t. Then, the optimality equations can be written as

$$Q_t^*(b_t) = \max\Bigg\{ r_t(b_t, DN) + \sum_{s_t \in \bar{\mathcal{S}}} b_t(s_t) \Bigg[\sum_{o_t \in O_{DN}} z_t(o_t|s_t, DN) \sum_{s_{t+1} \in \bar{\mathcal{S}}} p_t(s_{t+1}|s_t, DN, o_t) Q_{t+1}^*(b_{t+1}) \Bigg],$$
$$r_t(b_t, M) + \sum_{s_t \in \bar{\mathcal{S}}} b_t(s_t) \Bigg[z_t(M-|s_t, M) \sum_{s_{t+1} \in \bar{\mathcal{S}}} p_t(s_{t+1}|s_t, M, M-) Q_{t+1}^*(b_{t+1}) \Bigg]$$
$$+ b_t(0) z_t(M+|0, M) p_t(s_{t+1}|0, M, M+) Q_{t+1}^*(b_{t+1}) + \sum_{s_t \in \{1,2\}} b_t(s_t) z_t(M+|s_t, M) R_t(s_t) \Bigg\},$$
$$\forall t \in \mathcal{T},\ \forall b_t \in \mathcal{B}(\bar{\mathcal{S}}), \tag{4.2a}$$

$$Q_N^*(b) = R_N(b), \quad \forall b \in \mathcal{B}(\bar{\mathcal{S}}). \tag{4.2b}$$

We refer to the POMDP model formulated by the equations (4.2a) - (4.2b) as the *Unconstrained Partially Observable Breast Cancer Screening Model* (UPOBCS).

As the patient gets older, she may make a transition to another belief state as a result of the selected actions and corresponding observations. We use Bayesian updating to specify the transitions between belief states. More specifically, given that the patient has a belief state b_t at time t, takes action $a_t \in \mathcal{A}$ and observes $o_t \in O_{a_t}$, the patient's updated belief state $b_{t+1} \in \mathcal{B}(\bar{\mathcal{S}})$ is computed as follows:

$$b_{t+1}(s_{t+1}) = \frac{\sum_{s_t \in \bar{\mathcal{S}}} b_t(s_t)\ z_t(o_t|s_t, a_t)\ p_t(s_{t+1}|s_t, a_t, o_t)}{\sum_{s_t \in \bar{\mathcal{S}}} b_t(s_t)\ z_t(o_t|s_t, a_t)}, \quad \forall s_{t+1} \in \bar{\mathcal{S}}. \tag{4.3}$$

4.2.5 Constrained POMDP Model

We aim to obtain a mammography screening policy, $\pi \in \Pi$, that satisfies the constraints on the expected number of total mammography screenings, where Π represents the set of all policies. Let $c_t(b_t, a_t)$ denote the cost of choosing action a_t at time t when the patient's belief state is b_t

and C_0 denote the total number of available mammography screenings from decision epoch 0 onwards. We formulate the following constrained POMDP (CoPOMDP) model to explicitly account for mammography screening limit constraints:

$$\textbf{CoPOMDP: } \max_{\pi \in \Pi} E_{\delta}^{\pi}\Big[\sum_{t=0}^{N-1} r_t(b_t, a_t) + R_N(b_N)\Big] \tag{4.4a}$$

$$\text{s.t. } E_{\delta}^{\pi}\Big[\sum_{t=0}^{N-1} c_t(b_t, a_t)\Big] \leqslant C_0, \tag{4.4b}$$

where δ is the probability distribution for the patient's belief states at the first decision epoch, i.e., $\delta(b)$ satisfies $\sum_{b \in \mathcal{B}(\bar{\mathcal{S}})} \delta(b) = 1$ and $\delta(b) \geq 0,\ \forall b \in \mathcal{B}(\bar{\mathcal{S}})$. As the constraint (4.4b) restricts the expected number of total mammography screenings, we use $c_t(b_t, DN) = 0$ and $c_t(b_t, M) = 1$ as the cost values. We impose the limits on the expected number of total mammography screenings instead of the actual number of total exams as we would like to take into account all the possible scenarios for the patient in terms of her screening history.

Note that the limit on the number of mammography screenings is not a hard limit and the number of screenings recommended depends on the patient's screening history (i.e., the system history $\{(a_0, o_0), \ldots, (a_t, o_t)\},\ t \in \mathcal{T}$). More specifically, the patient has a stochastic process governing the screening outcomes and she might follow different paths (or scenarios), $\{(a_0, o_0), \ldots, (a_{N-1}, o_{N-1})\}$, throughout the decision process. As such, the number of mammography screenings recommended to the patient might be different for different paths. For instance, if the patient always makes negative observations as a result of the actions taken, then she may be recommended fewer screenings than the available number of screenings. On the other hand, the screening limit may be exceeded for the scenario where the patient makes multiple positive observations throughout the decision process. However, constraint (4.4b) ensures that the expected value of the number of screenings over these scenarios does not exceed the imposed limit.

The optimality equations given in (4.2a) - (4.2b) equivalently represent the objective function

(4.4a) and any solution algorithm for solving the optimality equations for POMDPs such as Monahan's algorithm [Monahan, 1982] and Sondik's algorithm [Smallwood and Sondik, 1973] can be used to find the optimal solution for the unconstrained problem formulated by only (4.4a). However, there is no efficient exact solution algorithm for the constrained POMDP model given by (4.4a) - (4.4b). We use grid-based approximations to solve CoPOMDP, which we describe in the next section.

4.3 Solution Methodology

Several studies in the literature that consider resource constraints in optimizing medical decisions convert all QALY values to monetary terms using a value for willingness-to-pay per QALY and focus on the resulting single-objective model [Erenay, 2010, Lee et al., 2009]. This approach is generally useful for reducing model complexity as it avoids modeling resource constraints directly. However, it is difficult to justify the use of willingness-to-pay (i.e., putting a dollar value for life years) from an ethical standpoint when the objective is to optimize the patients' well being [Neumann and Weinstein, 2010]. Furthermore, determination of a reasonable willingness-to-pay value is controversial even when the ethical issues are left aside. Therefore, we aim to explore the trade-off between health care costs and population health, and propose a modeling framework that explicitly incorporates the resource constraints.

Our solution methodology is based on a grid-based approximation for POMDP models. We first develop a grid-based approximation for UPOBCS where the mammography limit constraint is ignored. Then, we reformulate the approximation as a linear program (LP) and discuss how to augment the LP model with the mammography screening limit constraint. Finally, we introduce a set of constraints and binary variables to this model to ensure that we obtain deterministic and control-limit policies.

4.3.1 Grid-based Approximation for POMDPs

POMDP models are usually computationally intractable and state-of-the-art exact solution algorithms can only solve unconstrained problems with few core states to optimality. Therefore, several approximation methods are developed for POMDPs and grid-based approaches are frequently used methodologies to solve unconstrained POMDP models [Lovejoy, 1991b, Sandikci et al., 2013]. More specifically, we can find an approximate solution to the value functions (4.2a) - (4.2b) by discretizing the belief space into a finite set of grid points. Let $\mathcal{G} = \{g_1, \ldots, g_K\}$ be the grid set and $\mathcal{K} = \{1, \ldots, K\}$ be the index set of $\mathcal{G}$. Given that $\mathcal{G}$ includes all corner points of the belief simplex, any updated belief state $b_{t+1} \notin \mathcal{G}$ can be expressed as a convex combination of the points in $\mathcal{G}$, that is,

$$b_{t+1} = \sum_{k \in \mathcal{K}} \beta_k \boldsymbol{g}_k,$$

where $\sum_{k \in \mathcal{K}} \beta_k = 1$ and $\beta_k \geq 0$, for all $k \in \mathcal{K}$. Assuming that $\hat{Q}_{t+1}(\cdot)$ represents the approximate value functions, we obtain an approximation for the UPOBCS by replacing $Q^*_{t+1}(b_{t+1})$ with $\sum_{k \in \mathcal{K}} \beta_k \, \hat{Q}_{t+1}(\boldsymbol{g}_k)$ in the optimality equations (4.2a).

There may be more than one set of β_k values that span the updated belief state b_{t+1} and choosing a combination that yields better approximate values for $Q^*(b_{t+1})$ is important to find good approximations for the optimal value functions. Sandikci et al. [2013] show that, β_k values obtained by solving the following linear program give the tightest approximation for the optimal value functions for the given grid set $\mathcal{G}$:

$$\textbf{Beta\_M:} \min \left\{ \sum_{k \in \mathcal{K}} \beta_k \hat{Q}_{t+1}(\boldsymbol{g}_k) : \quad \sum_{k \in \mathcal{K}} \beta_k \boldsymbol{g}_k = b_{t+1}, \; \sum_{k \in \mathcal{K}} \beta_k = 1, \text{ and } \beta_k \geq 0 \; \forall k \in \mathcal{K} \right\}. \quad (4.5)$$

Observe that, to find the correct β-weights, this LP must be solved at every decision epoch, for each action, observation and grid point, as they all contribute to the value of the updated belief state. Therefore, we represent the interpolation weights with the vector $\beta_{tk}^{(\boldsymbol{g}_i, a_t, o_t)}$ which specifies the weight assigned to the kth grid point in $\mathcal{G}$, if the updated belief state is $\boldsymbol{g}_i$, the action taken is

a_t and the observation is o_t. As in Section 3.3.2, we can speed up the calculation of the β values by focusing on the *close* grid points, instead of calculating the β values iteratively during the backward induction. However, our state space and the grid set is much smaller than the ones in Chapter 3. Similarly, size of the LP (Beta_M) is also much smaller. As such, we carry out the computation of the β values using the iterative approach. In our base-case setting, obtaining the β-values takes approximately 30 minutes.

There are several approaches to construct the grid set $\mathcal{G}$ [Lovejoy, 1991b]. For instance, if we use a fixed-resolution uniform grid approach, total number of grid points is $|\mathcal{G}| = \binom{|\bar{\mathcal{S}}|+u-1}{\bar{\mathcal{S}}-1}$, where u is the resolution parameter that specifies the increments between grid points. We have three partially observable states, i.e., $|\bar{\mathcal{S}}| = 3$, and if we use a resolution value of two, we obtain the following six grid points:

$$[1,0,0],\ \ [1/2,1/2,0],\ \ [1/2,0,1/2],\ \ [0,1,0],\ \ [0,1/2,1/2],\ \ [0,0,1],$$

where the grid point $[1,0,0]$ corresponds to a perfectly healthy patient with no risks of in situ or invasive cancers. We generally need the resolution value to be as high as possible so that we obtain more accurate approximations for the POMDP model. However, increasing the resolution value may lead to very large grid sets that are computationally challenging to use in an approximation scheme. For instance, in a uniform grid, the resolution value should be at least 500 to differentiate the cancer risk between grid points by 0.2%, and we need 125,751 grid points to obtain such accuracy.

While the belief space $\mathcal{B}(\bar{\mathcal{S}})$ contains every possible belief state, the majority of those belief states are not reflective of realistic breast cancer risks. For instance, no patients would have a belief state corresponding to zero probability of being healthy, unless they are diagnosed with breast cancer. Therefore, we use a variable-resolution uniform grid approach to generate grid points, which helps us focus on specific parts of the belief space rather than generating equally-spaced grid points. For this purpose, we partition the belief space into n regions according to

the first dimension of the belief states. As the first dimension of a belief state corresponds to the probability of patient being healthy and the summation of other two dimensions corresponds to the probability of patient having cancer, this approach leads to a natural partitioning of the belief space. Let ρ_i represent the ith region in the belief space and u_i represent the resolution value used for region i. Then, the belief state b occupies ρ_i if $b(0) \in [\phi_{i-1}, \phi_i)$, where $0 \leq \phi_{i-1} < \phi_i \leq 1$. Thus, we construct the regions so that ρ_1 consists of the belief states that represent the patients with lowest breast cancer risk and ρ_n consists of the belief states that represent the patients with highest breast cancer risk.

Our numerical analysis indicates that the likelihood of a patient's belief state occupying a specific region decreases as we move away from ρ_1 to ρ_n. As a result, we use higher resolution values to sample grid points from ρ_1. Total number of grid points obtained using the variable resolution uniform grid approach can be calculated as follows:

$$|\mathcal{G}| = \binom{|\bar{\mathcal{S}}| + \lfloor (1-\phi_1)u_1 \rfloor - 1}{|\bar{\mathcal{S}}| - 1} + \sum_{i=2}^{n} \left[\binom{|\bar{\mathcal{S}}| + \lfloor (1-\phi_i)u_i \rfloor - 1}{|\bar{\mathcal{S}}| - 1} - \binom{|\bar{\mathcal{S}}| + \lfloor (1-\phi_{i-1})u_i \rfloor - 1}{|\bar{\mathcal{S}}| - 1} \right].$$

Figure 4.1 shows an example resolution setting with $n = 4$ regions, where the grid set $\mathcal{G}$ consists of 315 grid points. This resolution setting helps us separate the grid points by 0.002 in region one and the separation between grid points increases as the grid points approach region four.

After constructing the grid set $\mathcal{G}$ and calculating the convex combination weights, β, we can approximate the optimal value function of the POMDP model by formulating an MDP model over the grid set $\mathcal{G}$. Transition probabilities of the MDP model are $\mathbb{F}_t(\boldsymbol{g}_i|\boldsymbol{g}_k, a) = \sum_{o \in O_a} \mathcal{F}_t(\boldsymbol{g}_i|\boldsymbol{g}_k, a, o)$, where

$$\mathcal{F}_t(\boldsymbol{g}_i|\boldsymbol{g}_k, a, o) = \begin{cases} \beta_{ti}^{(\boldsymbol{g}_k, a, o)} \sum_{s_t \in \bar{\mathcal{S}}} g_k(s_t) z_t(o|s_t, a) \sum_{s_{t+1} \in \bar{\mathcal{S}}} p_t(s_{t+1}|s_t, a, o), & \text{if } a = DN,\ o \in O_{DN}, \\ & \text{or } a = M,\ o = M-, \\ \beta_{ti}^{(\boldsymbol{g}_k, a, o)} g_k(0) z_t(o|0, a) \sum_{s_{t+1} \in \bar{\mathcal{S}}} p_t(s_{t+1}|0, a, o), & \text{if } a = M \text{ and } o = M+. \end{cases}$$

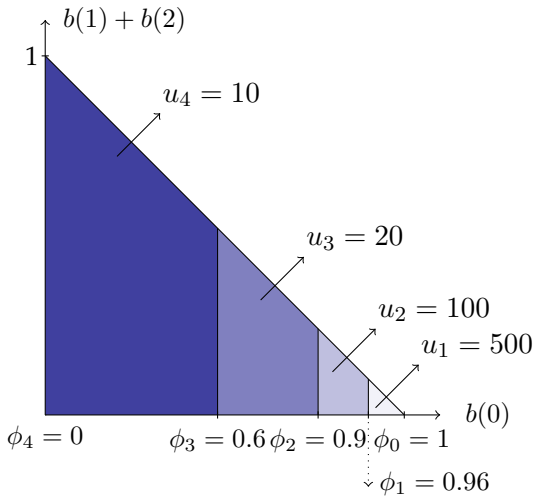

Figure 4.1: An example resolution setting with $\mathbf{n} = \mathbf{4}$ regions

A more detailed description of the MDP model is provided in Appendix B.1.

4.3.2 Linear Programming Model Formulation

We can use the following linear program to find an approximate solution for UPOBCS [Kallenberg and Cornelis, 1994]:

$$\textbf{BCS-LP: max} \quad h(x) \tag{4.7a}$$

$$\text{s.t.} \quad \sum_{a \in \mathcal{A}} x_0(\boldsymbol{g}_i, a) = \delta_i, \quad i \in \mathcal{K}, \tag{4.7b}$$

$$\sum_{a \in \mathcal{A}} x_t(\boldsymbol{g}_i, a) - \sum_{a \in \mathcal{A}} \sum_{k \in \mathcal{K}} \mathbb{F}_{t-1}(\boldsymbol{g}_i | \boldsymbol{g}_k, a) x_{t-1}(\boldsymbol{g}_k, a) = 0, \ \ i \in \mathcal{K}, \ 0 < t < N, \tag{4.7c}$$

$$x_N(\boldsymbol{g}_i) - \sum_{a \in \mathcal{A}} \sum_{k \in \mathcal{K}} \mathbb{F}_{N-1}(\boldsymbol{g}_i | \boldsymbol{g}_k, a) x_{N-1}(\boldsymbol{g}_k, a) = 0, \quad i \in \mathcal{K}, \tag{4.7d}$$

$$x_N(\boldsymbol{g}_i) \geq 0, \quad i \in \mathcal{K}, \quad x_t(\boldsymbol{g}_i, a) \geq 0, \quad a \in \mathcal{A}, \ i \in \mathcal{K}, \ t < N, \tag{4.7e}$$

where $h(x)$ corresponds to the expected total QALYs of a patient for a given solution x and is

defined as

$$h(x) = \sum_{t<N} \sum_{a\in\mathcal{A}} \sum_{k\in\mathcal{K}} r_t(\boldsymbol{g}_k, a) x_t(\boldsymbol{g}_k, a) + \sum_{t<N} \sum_{k\in\mathcal{K}} \sum_{s\in\{1,2\}} g_k(s) z_t(M+|s, M) R_t(s) x_t(\boldsymbol{g}_k, M) + \sum_{k\in\mathcal{K}} R_N(\boldsymbol{g}_k) x_N(\boldsymbol{g}_k).$$

The continuous variables $x_t(\boldsymbol{g}_k, a)$ correspond to the probability of occupying belief state $\boldsymbol{g}_k$ at time t and taking action a. Similarly, $x_N(\boldsymbol{g}_k)$ denotes the probability of occupying the belief state $\boldsymbol{g}_k$ at the final decision epoch. Also, $\mathbb{F}_t(\boldsymbol{g}_i|\boldsymbol{g}_k, a)$ specifies the transition probabilities at time t, given the belief state $\boldsymbol{g}_k$ and action a.

The objective function of BCS-LP maximizes the expected total QALYs of the patient. Constraints (4.7b) guarantee that the fraction of time a patient occupies belief state $\boldsymbol{g}_i$ at decision epoch 0 is equal to the initial probability value, δ_i, assigned to the corresponding belief state. Constraints (4.7c) and (4.7d) balance the flow in and out of each belief state at a given decision epoch. Finally, constraints (4.7e) state bound and logical restrictions on the variables.

The probability measure described by continuous x variables in BCS-LP provides a natural way to model the resource restrictions. We introduce the constraint

$$\sum_{t\in\mathcal{T}} \sum_{k\in\mathcal{K}} x_t(\boldsymbol{g}_k, M) \leq C_0 \tag{4.8}$$

to limit the expected number of total mammography screenings. We refer to the linear programming formulation obtained by combining BCS-LP and constraint (4.8) as CoBCS-LP.

4.3.2.1 Deterministic Mammography Screening Policies

One issue with CoBCS-LP is that its optimal solution does not necessarily correspond to a deterministic optimal policy (see Appendix B.2). Instead, we may obtain a randomized decision rule for any belief state $\boldsymbol{g}_k$, which is impractical in a clinical context. To overcome this issue, we

introduce binary variable $\theta_t(\boldsymbol{g}_k)$, which equals 1 if mammography screening is recommended for the patient in belief state $\boldsymbol{g}_k$ and decision epoch t and 0 otherwise. We add the following constraints to BCS-LP to ensure that the optimal policies are deterministic:

$$\begin{aligned} &x_t(\boldsymbol{g}_k, M) \leq \theta_t(\boldsymbol{g}_k), && \forall t \in \mathcal{T},\ \forall k \in \mathcal{K}, \\ &x_t(\boldsymbol{g}_k, DN) \leq 1 - \theta_t(\boldsymbol{g}_k), && \forall t \in \mathcal{T},\ \forall k \in \mathcal{K}, \\ &\theta_t(\boldsymbol{g}_k) \in \{0, 1\} && \forall t \in \mathcal{T},\ \forall k \in \mathcal{K}, \end{aligned} \tag{4.9}$$

where θ-variables guarantee that only one of the DN and M actions can be chosen for a belief state $\boldsymbol{g}_k$ at time t. Note that if $\theta_t(\boldsymbol{g}_k) = 1$, then $x_t(\boldsymbol{g}_k, M)$ can take any value between 0 and 1, and there could be multiple optimal solutions when both $x_t(\boldsymbol{g}_k, DN)$ and $x_t(\boldsymbol{g}_k, M)$ are equal to 0. However, we can disregard the values of x-variables to determine screening policies as θ-variables provide the resulting screening policy.

4.3.2.2 Control-limit Type Mammography Screening Policies

Similar to the deterministic policies, control-limit type policies are also desirable from a clinical perspective as they are easier for interpretation and implementation. A control-limit type policy is a policy for which there exists a belief state b^*, called the control-limit, such that one of the actions is prescribed for all belief states $b \leq b^*$ and the other action for all belief states $b > b^*$. We use first-order stochastic dominance for probability distributions to compare the belief states. That is, a belief state $\bar{b}$ is *stochastically larger* than the belief state b, denoted as $b \leq_s \bar{b}$, if $\sum_{i=j}^{|\bar{\mathcal{S}}|} b(i) \leq \sum_{i=j}^{|\bar{\mathcal{S}}|} \bar{b}(i)$ for any $j \in \{1, \ldots, |\bar{\mathcal{S}}|\}$. Therefore, we assume that the belief state $\bar{b}$ is *worse* than the belief state b, i.e., $\bar{b}$ corresponds to a patient with higher risk of developing breast cancer. We add the following constraints to the BCS-LP to obtain control-limit type policies:

$$\theta_t(\boldsymbol{g}_i) \leq \theta_t(\boldsymbol{g}_k), \quad t < N,\ i \in \mathcal{K},\ k \in \mathfrak{K}(i) = \{\hat{k} \in \mathcal{K} : \boldsymbol{g}_i \leq_s \boldsymbol{g}_{\hat{k}}\}. \tag{4.10}$$

Constraint (4.10) states that whenever the mammography action is chosen for the belief state $\boldsymbol{g}_i$, the same action should also be chosen for the belief states that are worse than $\boldsymbol{g}_k$. Note that the number of such constraints can be as large as $|N||\mathcal{G}|^2$. However, we can reduce the number of constraints needed to ensure control-limit type policies by using the stochastic ordering hierarchy between the grid points. For instance, if $\boldsymbol{g}_0 \leq_s \boldsymbol{g}_1 \leq_s \boldsymbol{g}_2$, we use following inequalities to ensure control-limit policies: $\theta(\boldsymbol{g}_0) \leq \theta(\boldsymbol{g}_1)$, $\theta(\boldsymbol{g}_0) \leq \theta(\boldsymbol{g}_2)$, and $\theta(\boldsymbol{g}_1) \leq \theta(\boldsymbol{g}_2)$. However, second inequality is redundant as it is already implied by first and third inequalities, and therefore need not to be included in our linear programming formulation.

We construct a directed graph, $\mathcal{D} = (V(\mathcal{D}), A(\mathcal{D}))$, that shows the stochastic ordering between grid points (see Figure 4.2). Nodes of the graph corresponds to grid points and there is an arc from $\boldsymbol{g}_i$ to $\boldsymbol{g}_k$ if $\boldsymbol{g}_i \leq_s \boldsymbol{g}_k$ and $\nexists\, \boldsymbol{g}_j$ such that $\boldsymbol{g}_i \leq_s \boldsymbol{g}_j \leq_s \boldsymbol{g}_k$ (see Appendix B.3 for an algorithm to construct $\mathcal{D}$). For example, in Figure 4.2, there is no arc between $\boldsymbol{g}_0$ and $\boldsymbol{g}_2$ as $\boldsymbol{g}_1$ satisfies the relation $\boldsymbol{g}_0 \leq_s \boldsymbol{g}_1 \leq_s \boldsymbol{g}_2$. Then, we can describe the control-limit policy constraints as follows:

$$\theta_t(\boldsymbol{g}_i) \leq \theta_t(\boldsymbol{g}_k), \quad t < N,\ (i,k) \in A(\mathcal{D}). \tag{4.11}$$

As such, we only need $O(N|G|)$ inequalities of type (4.11) compared to $O(N|G|^2)$ of type (4.10) to enforce control-limit policies.

We refer to the formulation obtained by combining the equations (4.7a) - (4.9), and (4.11) as Constrained Breast Cancer Screening Problem - Mixed-Integer Linear Programming formulation (CoBCS-MIP)

4.4 Mixed-Integer Programming Model For Resource Allocation

We extend CoBCS-MIP to devise a resource allocation scheme between patients with different risk characteristics. In this model, each subpopulation follows their own progression and patients share C_0 mammography screenings. We denote the set of patient types by $\mathcal{L}$ and each patient type

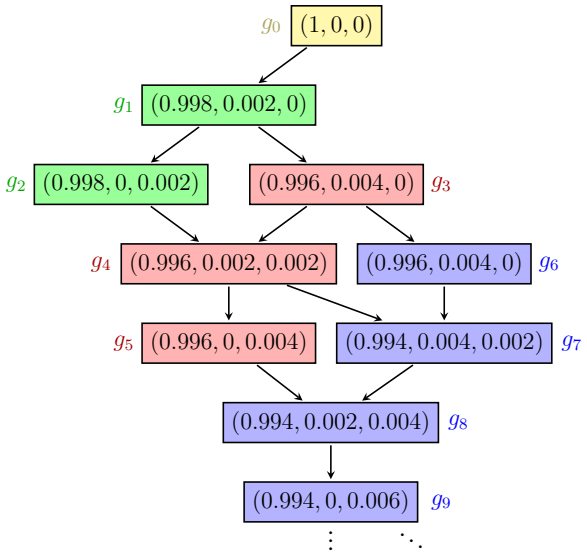

Figure 4.2: Graph of stochastic orderings between grid points

is assigned a weight of α_ℓ, representing the proportion of type ℓ patient in a given population, where $\sum_{\ell \in \mathcal{L}} \alpha_\ell = 1$. We formulate the resource allocation problem as follows:

$$\textbf{ExtCoBCS-MIP:} \quad \max \quad \sum_{\ell \in \mathcal{L}} \alpha_\ell h(x^\ell) \tag{4.12a}$$

$$\text{s.t.} \quad \sum_{\ell \in \mathcal{L}} \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{K}} \alpha_\ell x_t^\ell(\boldsymbol{g}_k, M) \leq C_0, \tag{4.12b}$$

$$(x^\ell, \theta^\ell) \in Y^\ell, \quad \forall \ell \in \mathcal{L}, \tag{4.12c}$$

where Y^ℓ is defined as

$$
\begin{aligned}
Y^\ell = \Bigg\{ & (x,\theta) : \sum_{a\in\mathcal{A}} x_0(\boldsymbol{g}_i,a) = \delta_i^\ell, \quad i \in \mathcal{K}, \\
& \sum_{a\in\mathcal{A}} x_t(\boldsymbol{g}_i,a) - \sum_{a\in\mathcal{A}}\sum_{k\in\mathcal{K}} \mathbb{F}_{t-1}^\ell(\boldsymbol{g}_i|\boldsymbol{g}_k,a)\, x_{t-1}(\boldsymbol{g}_k,a) = 0, \quad t < N,\ i \in \mathcal{K}, \\
& x_N(\boldsymbol{g}_i) - \sum_{a\in\mathcal{A}}\sum_{k\in\mathcal{K}} \mathbb{F}_{N-1}^\ell(\boldsymbol{g}_i|\boldsymbol{g}_k,a)\, x_{N-1}(\boldsymbol{g}_k,a) = 0, \quad i \in \mathcal{K}, \\
& \theta_t(\boldsymbol{g}_i) \le \theta_t(\boldsymbol{g}_k), \quad t < N,\ (i,k) \in A(\mathcal{D}), \\
& x_t(\boldsymbol{g}_i, M) \le \theta_t(\boldsymbol{g}_i), \quad t < N,\ i \in \mathcal{K}, \\
& x_t(\boldsymbol{g}_i, DN) \le 1 - \theta_t(\boldsymbol{g}_i), \quad t < N,\ i \in \mathcal{K}, \\
& \theta_t(\boldsymbol{g}_i) \in \{0,1\}, \quad t < N,\ i \in \mathcal{K}, \\
& x_N(\boldsymbol{g}_i) \ge 0, \quad i \in \mathcal{K}, \quad x_t(\boldsymbol{g}_i,a) \ge 0, \quad a \in \mathcal{A},\ i \in \mathcal{K},\ t < N \Bigg\}.
\end{aligned}
$$

The objective of the ExtCoBCS-MIP model maximizes the weighted QALYs of the patients and constraint (4.12b) ensures that the resource limits are satisfied. Note that, as we scale the x-variables using α values, C_0 corresponds to the per patient screening limit assigned to a given population. Finally, constraints (4.12c) enforce the condition that each subpopulation follow their own progression under a deterministic control-limit policy throughout the decision process.

4.5 Parameter Estimation

Our constrained POMDP models are informed by validated data sources. The validity of our model and policy recommendations are based on these data sources and our comparisons with the closely related studies in the literature (see Appendix B.4). In particular, we estimate the

following parameters to implement our method described in Sections 4.2 and 4.4:

- *Health state transition probabilities*. We use the University of Wisconsin Breast Cancer Simulation (UWBCS), a validated microsimulation model of breast cancer epidemiology to estimate the probability of making transitions between health states. The UWBCS is able to replicate breast cancer incidence and mortality rates observed in the U.S. population between 1975 and 2010. The UWBCS was developed as part of the National Cancer Institute (NCI)-funded Cancer Intervention and Surveillance Modeling Network (CISNET) and has been used by policy makers such as United States Preventive Services Task Force (USPSTF 2009). More information about UWBCS can be found elsewhere [Fryback et al., 2006].

- *Personal breast cancer risk*. We use a modified version of the Gail model, a validated personal breast cancer risk prediction tool, to estimate the breast cancer risk of in situ and invasive cancer at age 40 [Gail et al., 1999], which is available on the NCI's website (http://www.cancer.gov/bcrisktool/). The estimated risks of in situ and invasive cancers are used to calculate the initial belief states for the patients. In particular, at age 40, we estimate the invasive cancer risk as 0.2%, 0.3% and 1.2% for low-risk, average-risk and high-risk patients, respectively. The risks at later ages are calculated by using the Bayesian belief update formula given in Equation (4.3).

- *Performance of mammography screening and self-detection*. We obtain age-specific sensitivity and specificity of the mammography from Kerlikowske et al. [2000]. We use Barton et al. [1999]'s study and Baxter et al. [2001]'s study to estimate sensitivity and specificity of the clinical breast exams and breast self exam respectively. Finally, we obtain the proportion of clinical breast exam and breast self exam from Elmore et al. [2005] and Messina et al. [2004] to calculate the sensitivity and specificity of self-detection by a patient.

- *Post-cancer life expectancies*. We use the Surveillance, Epidemiology, and End Results (SEER) data [Siegel et al., 2014], the most comprehensive data source on cancer outcomes in the

U.S., and the method developed by Arias [2014] to estimate the age-specific post-cancer life expectancies for patients who are receiving cancer treatment.

- *Quality of life adjustments for reward function*. An important input for our model is the disutility of the patient towards various screening and diagnostic procedures associated with mammography. While we use the best available evidence to estimate the quality of life scores representing society's preferences, we recognize that these parameters are subjective and may change from one patient to another. Therefore, we conducted an extensive sensitivity analysis by using various disutility values. For our base case analysis, we use the following disutility values in our reward function:
  - 0.5 days for a negative mammogram [Mandelblatt et al., 1992];
  - 2 weeks for a true-positive (TP) mammogram [Velanovich, 1995];
  - 4 weeks for a false-positive (FP) mammogram, which is noted to be higher than TP mammogram [Earle et al., 2000].

  Our explicit modeling of imposing constraints on the number of mammography screenings also allows us to ignore the quality-of-life scores and focus only on the life-years associated with a particular policy under various mammography screening limits. Hence, in each of our numerical results, we include an experiment assuming that all disutilities are equal to zero and estimate the outcomes of various policies in terms of life years (LYs).

- *Distribution of breast cancer risk groups in the population.* Graubard et al. [2010] report that 7.5% (95% CI = 6.2% to 9%) of the U.S. population is in the high-risk group for breast cancer and recommended for chemo-preventive interventions. We estimate the distribution of low-risk and average-risk women in the U.S. population from Weisstock et al. [2013]. As a result, in our experiments with the ExtCoBCS-MIP model, we estimate the percentages of low-risk, average-risk, and high-risk patients as 51.5%, 41%, and 7.5%, respectively.

4.6 Numerical Results

In this section, we present empirical results on CoBCS-MIP and ExtCoBCS-MIP models. We first investigate the computational performance of the grid-based approximation in Section 4.6.1. Then, we analyze the effects of mammography screening limits at the individual patient level. We demonstrate the changes in screening recommendations under different mammography screening limits in Section 4.6.2 and discuss the cost-effectiveness of allocating additional mammography screenings in Section 4.6.3. Next, we consider a resource sharing mechanism between patients with different risk profiles. In Section 4.6.4, we present the results on efficient allocation of mammography screenings between patients from different risk groups. Finally, in Section 4.6.5, we compare the proposed dynamic screening policies to the population-based fixed-interval screening guidelines.

Implementation details. We implement our models in C++ and solve them using Cplex12.4 on a Linux workstation with 24, 2.67 GHz Intel Xeon CPUs and 132 GB memory. In our base case grid settings, we partition the belief space into four regions. More specifically, we take $\phi_0 = 1,\ \phi_1 = 0.96,\ \phi_2 = 0.9,\ \phi_3 = 0.7,$ and $\phi_4 = 0$. We use the resolution values $u_1 = 1000$, $u_2 = 200$, $u_3 = 25$, and $u_4 = 10$ to sample from regions one, two, three, and four, respectively, yielding 1083 grid points. This resolution setting implies that we can attain at most 0.001 accuracy between the grid points that span the belief space and the level of accuracy decreases as we move away from region one, which is the most focused region in the belief space. The solution times for CoBCS-MIP and ExtCoBCS-MIP, using default Cplex settings, vary significantly depending on the number of grid points and the constraints on mammography screenings. For instance, for the given grid resolution setting, CPU time varies between 2 hours to 20 hours for different mammography screening limits. Also note that solving CoBCS-LP is fairly easy compared to solving MIP formulations, as incorporations of the constraints (4.9)-(4.10) lead to a significant increase in model complexity. We observe that CoBCS-LP can be solved under four minutes for the base case resolution setting.

4.6.1 Analysis of algorithmic performance

We observe that the Gail model's breast cancer risk estimates always lead to the belief states that fall into region one. Therefore, we set the area of region one as large as possible and use the highest resolution values to sample from this region. We remark that while expanding region one increases the number of grid points significantly, our results remain insensitive to further expansion of this region. We make similar observations for the resolution values as well. In particular, although higher resolution values tend to yield more accurate results, we do not observe significant changes in our policy findings with further increases in the resolution values. Table 4.2 shows the QALY values obtained for average-risk (AR) and high-risk (HR) patients when different resolution values are used to construct the grid set. These results show that the QALY values converge as we use higher resolution values to sample from belief space. However, as it is the case for the resolution setting 2000_200_25_10, CoBCS-MIP becomes computationally intractable for larger grid sets.

		no limit		*mammo limit: 8*		*mammo limit: 4*	
resolution setting	$\|\mathcal{G}\|$	AR	HR	AR	HR	AR	HR
250_100_10_5	129	40.220	39.128	40.209	38.997	40.178	38.876
500_100_25_10	345	40.200	39.114	40.195	38.987	40.176	38.855
1000_200_25_10	1083	40.194	39.110	40.192	38.983	40.174	38.855
2000_200_25_10	3504	40.192	39.110	-	-	-	-

Table 4.2: Comparison of grid resolutions

We also compare the grid-based approximation with the Monahan's algorithm [Monahan, 1982] for optimally solving the unconstrained POMDP. We observe that the grid-based model (grBased) well approximates the optimal solution for the unconstrained breast cancer screening problem in terms of QALY values and average number of screenings recommended (see Table 4.3). To our knowledge, there is no efficient exact solution algorithm for the constrained POMDPs. Therefore, we can only assess the performance of the grid-based approximation for the unconstrained problem.

	AR		HR	
method	QALY	avg. # mammo.	QALY	avg. # mammo.
Monahan	40.191	10.19	39.110	22.50
grBased	40.194	10.19	39.110	22.81

Table 4.3: Comparison of Monahan's algorithm with grid-based approximation

Figures 4.3 and 4.4 show the optimal breast cancer screening policies obtained by Monahan's algorithm and the grid-based approximation, respectively. The dark-colored areas show the risk combinations for which it is optimal to undergo mammography screening, and the light-colored areas show those for which it is optimal to wait. These results indicate that the grid-based model successfully replicates the optimal policies. However, it is important to note that Monahan's algorithm generally leads to more precise policies in terms of the screening thresholds as the grid-based approximation can only differentiate the belief states by 0.001.

4.6.2 Allocation of Limited Mammography Resources to Different Age Groups

An important policy question for policy makers is that for a given cohort of women, what ages are optimal to allocate the limited mammography screening capacity. This question was also investigated by the CISNET models, which provided evidence for the most recent USPSTF policy recommendation in the U.S. [Mandelblatt et al., 2016]. We simulate a single patient's lifetime using the policy obtained by solving the CoBCS-MIP model and determine the mammography screening ages under different scenarios. In particular, we consider two different scenarios for an individual patient: Always make a negative observation as a result of an action taken in any period (*scenario-1*), and make positive observations every 10 years, starting at age 45, as a result of do nothing actions (*scenario-2*).

As it is also noted in Section 4.2.5, CoBCS-MIP implicitly models the possible scenarios for the patient. Note that the limit is imposed on the expected number of total mammography screenings, not on the actual number of exams. Therefore, policy obtained by CoBCS-MIP

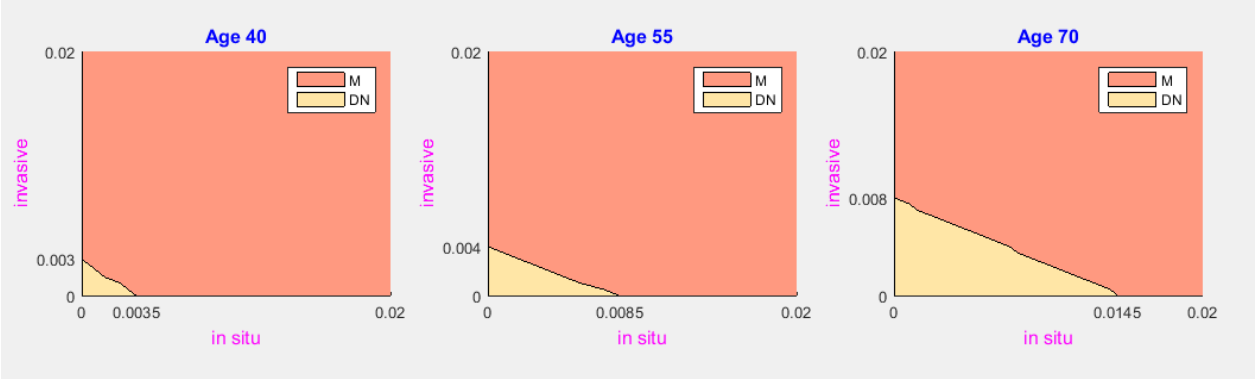

Figure 4.3: Optimal screening decisions obtained by Monahan's algorithm for different ages

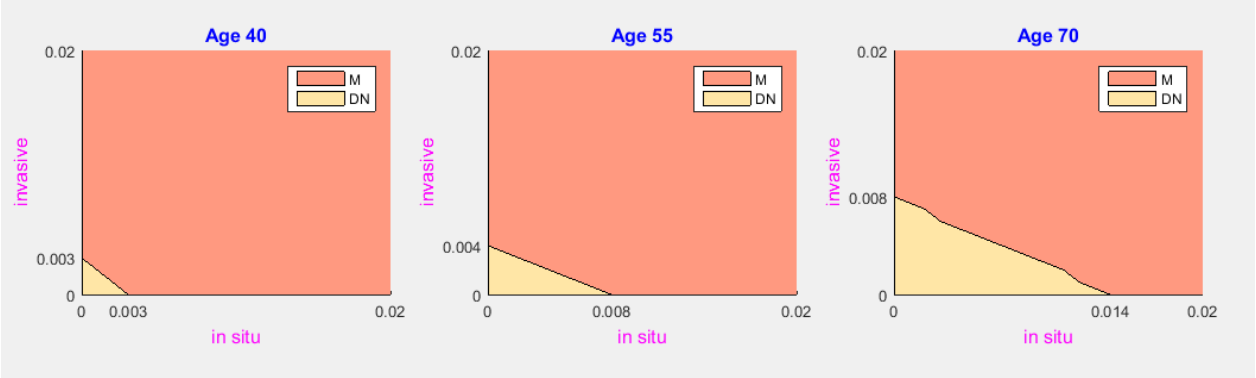

Figure 4.4: Screening decisions obtained by grid-based approximation for different ages

might yield number of mammography screenings exceeding the imposed limits when evaluated under a particular scenario. Such scenarios usually include do nothing actions with positive observations, where these observations elevate the breast cancer risk of the patient significantly. Figure 4.5 shows how the breast cancer risk of an average-risk patient changes until age 80 when she does not receive any screening. This figure shows that whenever the patient makes a positive observation, it leads to a significant increase in the patient's breast cancer risk and the magnitude of the increase depends on the patient's age. As such, the patient is usually recommended less mammography screenings than the imposed limits when she does not make positive observations and more mammography screenings are reserved for the cases in which

she makes positive observations. This behavior is more prevalently observed when the available mammography screenings are more limited.

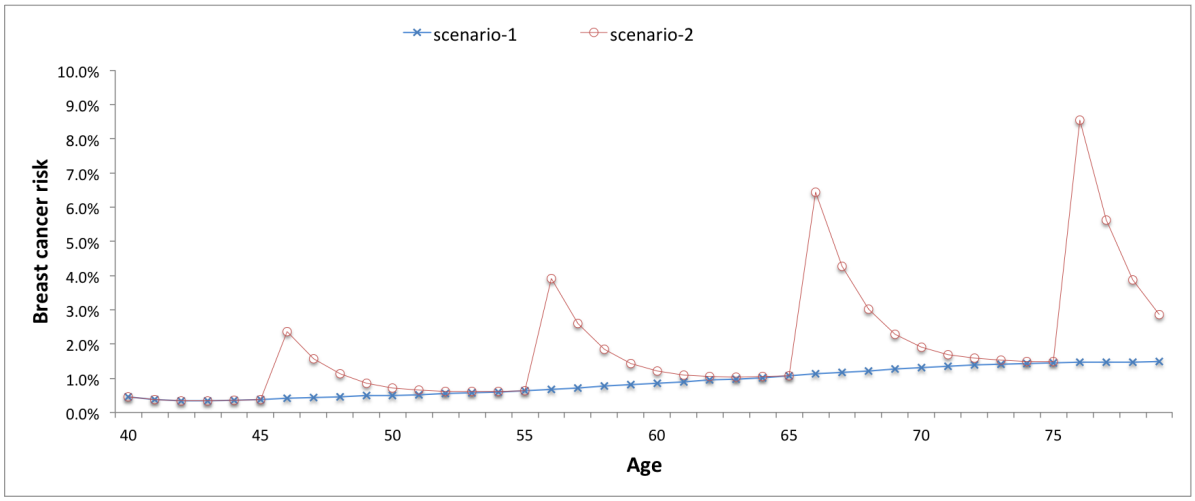

Figure 4.5: Evolution of the breast cancer risk for an average-risk patient under no screening case

We next investigate the policy recommendations for scenario-1. Table 4.4a provides results on how this patient's in situ and invasive cancer risks change over time and when she is recommended to undergo screening under different mammography screening limits. When there is no limit on the number of screenings, the patient undergoes mammography screenings starting at age 40 and the total number of mammography screenings is nine. The disutility of mammography and the risk of getting false-positive mammograms prevent the patient from undergoing more mammography screenings. If the mammography screening limit is eight, then this patient is assigned six mammography screenings throughout her life time. Similar trends are observed for other mammography screening limits as well. In particular, if the mammography screening limit is tight, (e.g., four) then the patient is not recommended any mammography screenings throughout her lifetime given that she always makes negative observations.

Table 4.4b presents the results for scenario-2, where the patient makes positive observations every 10 years, starting at age 45. As the breast cancer risk for the patient significantly increases with positive observations, she is usually recommended a screening at the next period. For

no limit				*mammo limit:8*				*mammo limit:4*			
age	insitu risk	invasive risk	action	age	insitu risk	invasive risk	action	age	insitu risk	invasive risk	action
40	0.0016	0.003	M	40	0.0016	0.003	M	40	0.0016	0.003	DN
48	0.0027	0.0017	M	50.5	0.0032	0.002	M	80	0.0097	0.0052	DN
53	0.0033	0.002	M	58	0.0048	0.0028	M				
57	0.0042	0.0025	M	62	0.0055	0.0031	M				
60.5	0.0049	0.0028	M	66.5	0.0067	0.0038	M				
64	0.0058	0.0032	M	71.5	0.0085	0.0046	M				
67	0.0063	0.0034	M	80	0.0096	0.005	DN				
70.5	0.0076	0.004	M								
74	0.0081	0.0043	M								
80	0.0094	0.0049	DN								
# of mammo screenings:			9				6				0

(a) All-negative observations

no limit				*mammo limit:8*				*mammo limit:4*			
age	insitu risk	invasive risk	action	age	insitu risk	invasive risk	action	age	insitu risk	invasive risk	action
40	0.0016	0.003	M	40	0.0016	0.003	M	40	0.0016	0.003	DN
45.5	0.0124	0.0088	M	45.5	0.0124	0.0088	M	45.5	0.0141	0.01	M
46	0.0037	0.0031	M	46.5	0.0035	0.0028	M	46.5	0.0038	0.0032	M
49	0.0027	0.0018	M	55.5	0.0227	0.0148	M	55.5	0.0227	0.0148	M
54.5	0.0037	0.0022	M	56	0.0059	0.0038	M	56	0.0059	0.0038	M
55.5	0.0112	0.0066	M	60.5	0.0052	0.003	M	65.5	0.0382	0.0235	M
57.5	0.0042	0.0025	M	65	0.0064	0.0036	M	66	0.007	0.0046	M
60.5	0.0047	0.0027	M	70	0.0079	0.0044	M	75.5	0.0526	0.0299	M
64	0.0058	0.0032	M	75	0.0088	0.0048	M	80	0.0097	0.0053	DN
65.5	0.0218	0.0122	M	80	0.0091	0.0046	DN				
67.5	0.0065	0.0036	M								
71	0.0077	0.0041	M								
75	0.0085	0.0045	M								
80	0.0091	0.0046	DN								
# of mammo screenings:			13				9				7

(b) Positive observations due to self-detection once in every 10 years starting at age 45

Table 4.4: Mammography allocation ages for a patient with initial belief distribution $[\mathbf{0.994}, \mathbf{0.0016}, \mathbf{0.003}]$ (i.e., she has a $\mathbf{0.16}\%$ risk of in situ cancer and $\mathbf{0.3}\%$ risk of invasive cancer)

instance, when there are no limits on available number of screenings, while the patient's in situ and invasive cancer risks are 0.37% and 0.22% at age 54.5, these values increase to 1.12% and 0.66%, respectively, after she makes a positive observation at age 55. As a result, the patient is assigned a mammography screening at age 55.5. Note that the patient might be recommended even more screenings than the imposed limits. In particular, when the limit is four, she is recommended seven mammography screenings, and these screenings come after the positive

observations (at age 45, 55, 65, and 75), except for the screening at age 40.

We next simulate one million average-risk patients and another one million high-risk patients and observe the screening allocation patterns for different ages to assess what ages are allocated the most mammography exams. Each patient enters the simulation at age 40 with an initial belief state selected according to the initial risk distribution of the patients and leaves the simulation either when she is diagnosed with cancer or when she dies. Figure 4.6 shows the average number of mammography screenings assigned to the patients between ages 40 and 80. Note that we report the average number of mammography screenings for two-year intervals, so the maximum number of mammography screenings recommended in this interval can be at most four, considering the 6-month decision epochs.

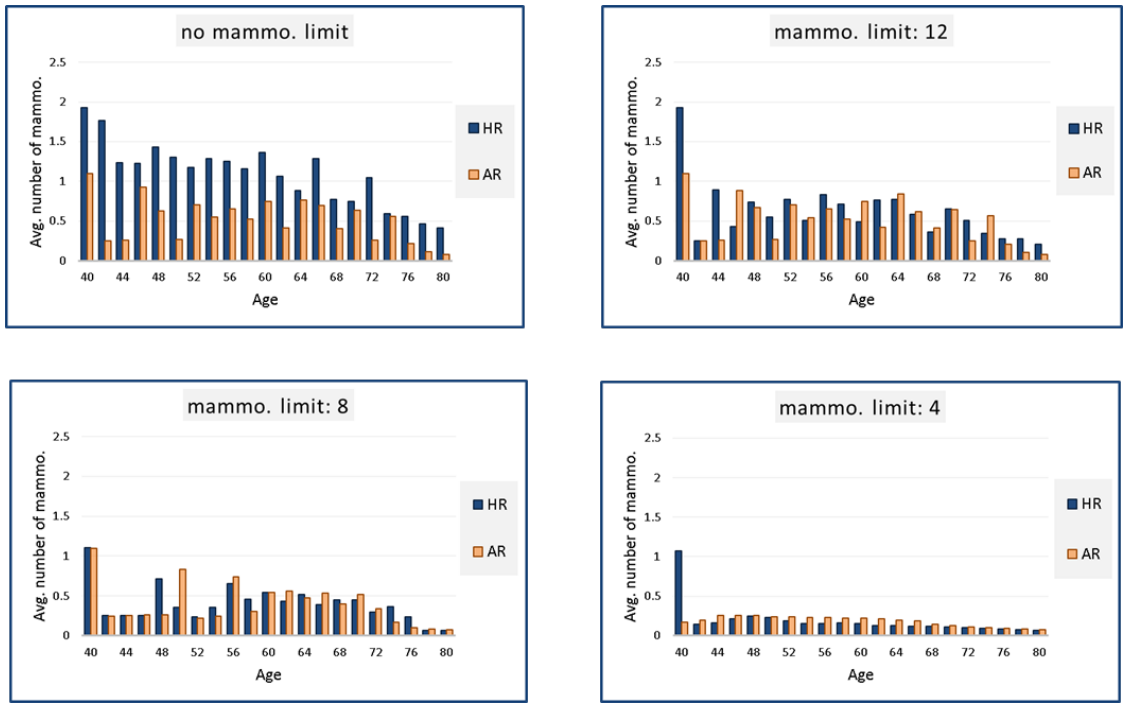

Figure 4.6: Average number of mammography screenings received by patients at different ages and risk levels

The simulation results show that, even when there is a limited number of available screenings, high-risk patients tend to receive screenings at earlier ages. On the other hand, immediate ages

for eliminating mammography screenings for average-risk patients are the ages between 40 and 50. More specifically, when there are no limits on the number of screenings, patients are allocated mammography screenings with no discernible pattern. However, if the available number of mammography screenings is eight per patient, then all the patients are recommended to undergo mammography screenings at age 40 and only patients making positive observations due to self-detection are recommended screenings until age 50. Note that if the mammography screening limit is four, only high-risk patients are recommended mammography screenings at age 40.

4.6.3 Effects of Mammography Screening Limit Constraints on Total Life-years and QALYs

By solving the CoBCS-MIP for a sequence of mammography screening limits, we determine incremental QALY gains attributable to consecutive levels of resources. Table 4.5 shows average QALY (avg. QALY) values obtained for different mammography screening limits as well as the QALY gains (gain (day)) per average-risk and high-risk patients. We repeat this experiment for different disutility values, namely, "base-case disutility", "low disutility" and "no disutility". We observe that although the first four screenings increase the QALY values considerably for average-risk patients, QALY gains become smaller for the subsequent mammography screenings. However, as the disutilities associated with mammography screenings decrease, QALY gains due to the additional screenings increase. On the other hand, high-risk patients benefit from increased number of screenings even when the highest disutility values are assumed for the screening.

We also conduct a cost-effectiveness analysis to demonstrate the benefits of additional mammography screenings allocated to the patients, summarized by the "Cost/QALY" columns in Table 4.5. We use the estimates reported in a recent study by Sprague et al. [2014] for various costs of mammography, false-positives, and cancer treatment. More specifically, Sprague et al. [2014]

	base case disutility			*low disutility*			*no disutility*		
limit	avg. QALY	gain (day)	Cost/QALY	avg. QALY	gain (day)	Cost/QALY	avg. LY	gain (day)	Cost/LY
0	39.430	-	-	39.430	-	-	39.430	-	-
2	40.123	253.4	$2,968	40.128	255.3	$2,944	40.146	261.6	$2,879
4	40.174	18.6	$10,414	40.183	20.1	$9,290	40.215	25.3	$7,423
6	40.185	3.7	$32,109	40.200	6.1	$22,360	40.245	11.2	$11,969
8	40.192	2.9	$50,389	40.212	4.2	$30,432	40.271	9.3	$13,835
10	40.194	0.7	$175,136	40.219	2.8	$45,514	40.292	7.8	$16,388
12	40.194	0.0	-	40.224	1.6	$81,327	40.310	6.6	$19,371
16	40.194	0.0	-	40.225	0.5	$271,975	40.338	10.2	$24,753
20	40.194	0.0	-	40.225	0.0	-	40.358	7.5	$33,109
24	40.194	0.0	-	40.225	0.0	-	40.374	5.8	$43,467

(a) Average-risk patient

	base case disutility			*low disutility*			*no disutility*		
limit	avg. QALY	gain (day)	Cost/QALY	avg. QALY	gain (day)	Cost/QALY	avg. LY	gain (day)	Cost/LY
0	36.677	-	-	36.677	-	-	36.677	-	-
2	38.667	727.3	$2,432	38.671	728.9	$2,427	38.694	737.2	$2,411
4	38.855	68.9	$4,119	38.864	70.5	$4,034	38.902	76.0	$3,669
6	38.930	27.4	$6,393	38.945	29.7	$5,881	38.997	34.8	$5,008
8	38.983	19.4	$7,904	39.003	21.1	$7,165	39.069	26.2	$5,699
10	39.022	14.1	$10,042	39.048	16.3	$8,809	39.127	21.4	$6,735
12	39.051	10.7	$12,904	39.083	12.7	$10,783	39.176	17.8	$7,682
16	39.090	14.2	$19,029	39.132	18.2	$14,891	39.253	28.3	$9,572
20	39.107	6.3	$41,929	39.160	10.2	$25,631	39.309	20.3	$12,905
24	39.110	1.1	$160,401	39.174	5.0	$51,889	39.350	15.0	$17,262

(b) High-risk patient

Table 4.5: QALY values and cost-effectiveness analysis for different mammography screening limits. Base case disutility: FP-disutility = 4 weeks, TP-disutility = 2 weeks, negative mammography disutility = 0.5 days. Low disutility: FP-disutility = 3 weeks, TP-disutility = 2 weeks, and negative mammography disutility = 0.25 days

report that the cost of mammography is $140, cost of diagnostic screening after mammography is $105, and cost of biopsy is $940. Also, treatment costs for patients with in situ and invasive cancers are reported as $12,000 and $22,000, respectively. Although QALY gains decrease as the patients are allocated more screenings, these additional mammography screenings may still be cost-effective. For example, for the average-risk patient, limiting the number of mammography screenings to six per patient has an incremental cost per QALY value of $\$32,109$ (over a mammography limit of four per patient) for the base case disutility values and $\$22,360$ for the low disutility values indicating that allowing two more mammography screenings per patient over

four mammography screenings is cost-effective (based on the well-accepted cost-effectiveness threshold of $50,000/QALY [Drummond, 2005]).

4.6.4 Allocation of Mammography Resources Among Patients from Different Risk Groups

We next estimate the allocation of limited mammography screenings to the patients with different risk levels using the ExtCoBCS-MIP model. We use the proportion of women in each risk group as given in Section 4.5. Table 4.6 provides the average number of mammography screenings assigned to 40-year-old patients from different risk groups, namely, low-risk (LR), average-risk (AR) and high-risk (HR), for different mammography screening limits. We also present the QALY values for each risk group and compare these QALY values with the case where each risk group is allocated equal number of mammography screenings. For instance, if the available number of screenings is four per patient (e.g., there are four million available mammography screenings when there are one million patients in the population), then LR, AR, and HR patients are allocated 2.88, 3.69, and 14.05 screenings, respectively. Note that we consider a per patient screening limit assigned to a given population and weighted average of the number of screenings assigned to each risk group with the ratio of each risk group in the population satisfies the imposed screening limit.

	Number of mammo.			*QALY by risk group (years)*			
mammo. limit	LR	AR	HR	LR	AR	HR	avg. QALY
4	2.88	3.69	14.05	40.389 (-2.1 d)	40.170 (-1.5 d)	39.077 (+80.1 d)	40.201 (+4.3 d)
6	3.73	6.61	18.31	40.394 (-2.0 d)	40.187 (+0.9 d)	39.102 (+62.7 d)	40.212 (+4.0 d)
8	5.47	8.89	20.58	40.398 (-0.7 d)	40.194 (+0.6 d)	39.108 (+45.7 d)	40.218 (+3.3 d)
10	7.53	10.77	22.81	40.400 (+0.0 d)	40.194 (+0.0 d)	39.109 (+31.9 d)	40.219 (+2.4 d)
no limit	7.56	10.82	23.00	40.400 (+0.0 d)	40.194 (+0.0 d)	39.110 (+0.0 d)	40.219 (+0.0 d)

Table 4.6: Mammography allocation between patients with different risk levels

Compared to the case where each risk group receives four mammography screenings, our resource allocation scheme leads to 2.1 and 1.5 days of QALY loss for LR and HR patients,

respectively, and 81 days of QALY gain for HR patients. Overall, the QALY gain for the population is 4.3 days. Note that, because the proportion of HR patients is much lower compared to other risk groups, it is possible to achieve significant QALY gains for HR patients by efficient distribution of the mammography screenings even when the screening resources are very limited.

The proportion of the mammography screenings assigned to the high-risk patients increases as the imposed limits become tighter. Figure 4.7 shows that while high-risk patients are allocated 16.9% of available resources when there are 10 screenings available per patient, they are allocated 26% of the resources when the limit is four. These results show that even when available resources are very limited, low-risk patients should still be allocated a non-negligible portion of the screening capacity.

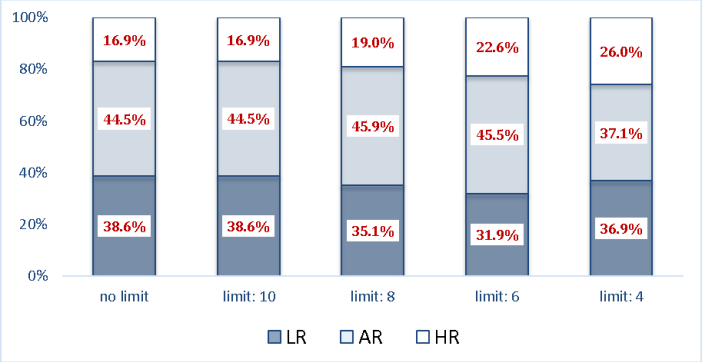

Figure 4.7: Percentage of mammography screenings allocated to each risk group for different screening limits

4.6.5 Comparison to the Population-based Screening Guidelines

As our last numerical experiment, we compare the dynamic screening policies obtained by solving ExtCoBCS-MIP with the fixed-interval screening policies recommended by several agencies such as USPSTF and ACR. Table 4.7 compares the performance of these population-based screening guidelines to that of our proposed policies when the same number of per patient mammography screenings is allowed (ECPOMDP-n represents the case where the mammography limit is n per

patient). For each policy, we present the average QALYs (avg. QALY), QALY gains over the no screening policy (gain (day)) and average number of mammography screenings (num. mammo.), and we repeat this experiment for different disutility values. Note that patients generally do not utilize all the mammography screenings that are assigned in fixed intervals as some patients are diagnosed with cancer at earlier ages or die before reaching age 90.

	base case disutility			*low disutility*			*no disutility*		
policy	avg. QALY	gain (day)	num. mammo.	avg. QALY	gain (day)	num. mammo.	avg. LY	gain (day)	num. mammo.
no screening	39.456	-	0	39.456	-	0	39.456	-	0
50-70 triennial	39.975	189.4	6.25	39.992	195.5	6.25	40.036	211.8	6.25
ECPOMDP-7	40.216	277.4	7.00	40.235	284.1	7.00	40.286	303.0	7.00
50-69 biennial	39.985	193.1	8.90	40.009	201.6	8.90	40.071	224.5	8.90
ECPOMDP-10	40.219	278.5	10.00	40.247	288.5	10.00	40.319	315.1	10.00
40-74 biennial	40.131	246.4	15.94	40.174	262.0	15.94	40.287	303.1	15.94
ECPOMDP-18	40.219	278.5	10.05	40.252	290.3	13.76	40.376	335.6	18.00
50-69 annual	39.959	183.6	17.62	40.007	201.0	17.62	40.129	245.4	17.62
ECPOMDP-20	40.219	278.5	10.05	40.252	290.3	13.76	40.385	339.1	20.00
40-74 annual	40.049	216.4	30.95	40.134	247.4	30.95	40.350	326.1	30.95
ECPOMDP-35	40.219	278.5	10.05	40.252	290.3	13.76	40.427	354.4	35.00

Table 4.7: Comparison of dynamic screening strategies with population-based screening guidelines. Base case disutility: FP-disutility = 4 weeks, TP-disutility = 2 weeks, negative mammography disutility = 0.5 days. Low disutility: FP-disutility = 3 weeks, TP-disutility = 2 weeks, and negative mammography disutility = 0.25 days

These results indicate that our proposed dynamic screening policies provide significant QALY gains over the population-based screening guidelines. For example, for the base-case disutility values, patients gain 85.4 days of more QALYs on average (39.985 years vs. 40.219 years) compared to 50-69 biennial screening, when they undergo 10 mammography screenings that are dynamically allocated. As the disutility values associated with mammography screening decrease, QALY gains of dynamic screening policies over the population-based screening guidelines decrease as well. However, an increase in the number of mammography screenings still has adverse effects on the patients in the form of unnecessary mammography exams and false-positive mammograms. In summary, this experiment shows that using our modeling approach to allocate limited resources instead of population-based screening policies has significant benefits in terms of both QALYs and LYs.

4.7 Concluding Remarks

We consider the problem of allocating limited breast cancer screening capacity to a target population to maximize the benefit of screening. For this purpose, we suggest a constrained POMDP model and we develop a variable-resolution grid-based approximation method to convert the constrained POMDP into a mixed-integer linear program. We introduce additional constraints to the mixed-integer linear programming formulation to ensure that the optimal policies are practical and clinically acceptable (i.e., deterministic and control-limit type). We use breast cancer epidemiology data from the U.S. to test how the allocation of screening capacity is affected as the mammography limit changes.

Our work could help to plan for mammography screenings in developing countries with limited screening resources as well as in various not-for-profit programs in developed countries that are targeting low-income population and are faced with budget restrictions. By using the proposed modeling scheme, policymakers could evaluate QALY and resource implications of different screening policies and make resource allocation decisions to maximize the effect of screening. According to the WHO, only a few countries have established breast cancer screening guidelines. Interestingly, even the least aggressive fixed-interval screening policy that is used in practice by developed countries (50-70 triennial screening) requires seven mammography screenings per patient, which may discourage other countries that are planning to implement breast cancer screening programs. In that regard, our findings suggest that allocating even a few mammography screenings per patient could result in significant QALY gains and such countries should consider implementing breast cancer screening programs based on their screening capacity.

One major benefit of our modeling approach over traditional POMDP applications in health care is the ability to explicitly model various constraints while optimizing a health outcome such as QALYs and LYs. In fact, many researchers and policymakers prefer optimizing over LYs as opposed to the QALYs, since the estimation of disutility values associated with screening

tests and treatments is very controversial. On the other hand, most studies that aim to optimize health outcomes need to use QALYs to consider the tradeoff between LYs and disutility values. For instance, if we ignore the various disutilities associated with mammography screening, the optimal policy will always recommend screening every 6 months for the unconstrained breast cancer screening problem. However, our modeling framework allows evaluating different mammography screening limits to find the one that balances different aspects of the mammography screenings such as LY gains and risk of getting a false-positive mammogram. Thus, our study eliminates the model dependence on QALYs and allows one to focus only on LYs as demonstrated in Section 4.6.

This work comes with some limitations. First, although the grid-based solution we offer is the best solution that can be obtained for a given set of grid points, we cannot assess the quality of this solution compared with the true optimum as there is no efficient exact solution methods for constrained POMDPs. However, considering that grid-based solution approach successfully approximates the optimal solution generated by Monahan's algorithm for the unconstrained POMDP model [Monahan, 1982], we believe that our solutions are not far from the true optimum. Second, we only consider three types of patients in a population especially due to limited data and computational resources. We acknowledge that finer categorization of the patients could lead to improved policy recommendations. Finally, we use the U.S.-specific data in our numerical experiments and recommendations for other countries may require country-specific breast cancer data. On the other hand, we expect our main policy insights such as high-risk patients should receive mammography at age 40 even when the resources are very limited to be valid in general.

5 USING SIMULATION MODELING TO INVESTIGATE THE IMPACT OF RACIAL DISPARITIES ON BREAST CANCER MORTALITY

5.1 Problem definition

In U.S., black women have lower breast cancer incidence than white women. However, breast cancer mortality rates have been consistently higher among black women and the disparity in breast cancer mortality between these race groups grows over the years. In particular, the breast cancer mortality among white women steadily decreased from 1990 onward with an average annual rate of 2.4%, while the rate of decrease among black women was only 1.1% over the same period [Smigal et al., 2006]. Figure 5.1 shows U.S.-observed breast cancer incidence and mortality trends for each race group.

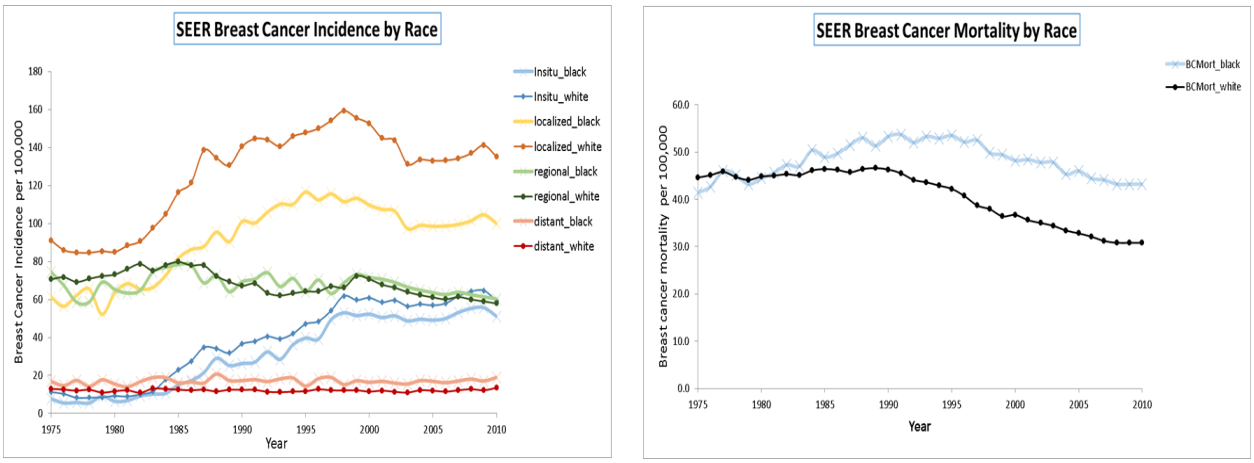

Figure 5.1: Comparison of SEER incidence and mortality rates between black and white women

There are several factors contributing to the racial disparities between black and white women. In particular, differences in mammography screening uptake, mammography effectiveness and treatment utilization have long been considered as the main factors contributing to the racial disparities. In addition, differences in disease natural history between black and white women

have also been recognized as an important factor.

Many cancer simulation studies model the disease natural history using natural history parameters. These parameters are referred as the unobservable parameters in a simulation model (e.g., tumor growth rate and ratio of benign tumors) that cannot be informed by the general data sources. In a recent study, Batina et al. [2013] use University of Wisconsin Breast Cancer Simulation Model (UWBCS) to evaluate the contribution of racial differences in tumor natural history to observed disparities in breast cancer incidence and find that mean tumor growth rate is 63.6% higher, and percentage of highly aggressive tumors are 2.2 times greater for black women.

In this study, we extend Batina et al. [2013]'s study to investigate the contributing factors for the differences in breast cancer mortality between black and white women. In our analysis, we use UWBCS of the Cancer Intervention and Surveillance Modelling Network (CISNET) to replicate U.S.-observed breast cancer incidence and mortality rates for both races. As UWBCS includes many natural history parameters, a lengthy calibration process is required to generate race-specific breast cancer simulation models. To this end, we propose a novel approach based on machine learning methods to speed up the calibration process and provide detailed numerical experiments to demonstrate the efficiency of our calibration method.

5.2 Methods

5.2.1 Overview

We modify the original UWBCS to obtain race-specific breast cancer simulation models for black and white women. Although we generally maintain the original model structure, we induce race-specific inputs for key components such as mammography screening uptake and treatment dissemination. The critical part of UWBCS in replicating observed breast cancer outcomes is the calibration of natural history parameters, which generally requires evaluation of millions of

parameter combinations via simulation. In this study, we propose a novel approach for speeding up the calibration process using machine learning techniques. We perform the calibration separately for black and white women and estimate the best fitting natural history parameters by calibrating to the observed 1975-2010 race-specific incidence data from the Surveillance, Epidemiology and End Results (SEER) program of the National Cancer Institute.

5.2.2 UWBCS model overview

UWBCS is a discrete-event, stochastic simulation model of breast cancer epidemiology in the U.S. female population and has been used to address a variety of breast cancer screening policy questions. The model is designed to match age and stage-specific breast cancer incidence rates and age-specific mortality rates in the U.S. female population between 1975 and 2010. The UWBCS incorporates observed age-specific U.S. screening patterns, secular trends in cancer risk, and dissemination of adjuvant treatment. The model outputs age-specific breast cancer incidence rates by stage and age-specific breast cancer mortality rates.

UWBCS has four core components: breast cancer natural history, breast cancer detection, breast cancer treatment and non-breast-cancer mortality among the U.S. women. While most of the parameters such as those used in treatment and mortality derived directly from the published studies, several parameters such as those used in natural history are unobservable and cannot be directly estimated from available data sources. Therefore, the UWBCS calibrates these natural history parameters to replicate the observed breast cancer incidence and mortality. More information about UWBCS can be found elsewhere [Fryback et al., 2006].

5.2.3 Race-specific adjustments in model parameters

In order to obtain race-specific breast cancer simulation models, namely, UWBCS-black for black women population and UWBCS-white for white women population, we modify the following input parameters of UWBCS:

- *Population age structure.* We use the race-specific population age structure from Census data [Census, 2016] to modify the population age structure of UWBCS to account for race.

- *Non-breast cancer mortality.* We estimate the cohort specific mortality from non-breast cancer causes using Berkeley Mortality Database [Berkeley, 2016] and the National Center for Health Statistics [CDC, 2016].

- *Mammography dissemination.* We use race-specific screening data from Breast Cancer Surveillance Consortium (BCSC) to generate race-specific mammography dissemination parameters [Cronin et al., 2009].

- *Breast cancer incidence.* We estimate seperate breast cancer incidence rates in the absence of screening for black and white women. More specifically, we adjust Holford et al. [2006]'s age-period-cohort (APC) model (the statistical model used by UWBCS to generate breast cancer incidence rates) for each race by using SEER-reported race-, sex- and age-specific breast cancer incidence rates.

- *ER/HER2 proportions.* Treatment type recommended to the patients are greatly affected by the tumor characteristics. Some tumor cells grow with the help of estrogen and progestrone hormones. Estrogen-receptor (ER) positive or progestrone-receptor (PR) positive tumors are more likely to respond to hormone therapy. Human epidermal growth factor receptor 2 (HER2) proteins are also important in cell growth and lead to aggressive tumor growth. On the other hand, HER2-positive tumors respond to anti-HER2 drugs such as Herceptin, and therefore have higher cure rates. Parise and Caggiano [2014] report that the ratio of triple negative tumors (ER-, PR-, HER2-) is significantly higher in black women compared to white women (25.4% vs 11.4%). As UWBCS does not consider PR-status, we only adjust the ER and HER2 proportions for each race using Parise and Caggiano [2014].

- *Treatment dissemination.* We use Mariotto et al. [2002, 2006] to estimate the use of adjuvant therapy for each race. Overall, we note that black women are 22% less likely to receive

multiagent chemotherapy and 15% less likely to receive hormonal therapy compared to white women.

Apart from the race-specific inputs, several other parameters are assumed to be identical for each race. In particular, we assume that chemotherapy and hormonal therapies are equally effective for black and white women [Newman, 2005, Vicini et al., 2010]. Also, sensitivity of mammography screening is assumed to be equal for each race group [Yankaskas and Gill, 2005].

5.2.4 Model calibration for natural history parameters

We calibrate the eight of ten natural history parameters of the original UWBCS to determine the race-specific values for which UWBCS successfully replicate U.S.-observed incidence and mortality rates. Table 5.1 summarizes these natural history parameters as well as the range of values assigned to each parameter during calibration. Among the natural history parameters, maximum size of in situ tumors and maximum size of limited malignant potential (LMP) tumors does not change by race. However, several other parameters such as proportion of LMP tumors, maximum sojourn time for LMP tumors after which LMP tumors regress and APC model related parameters (Onset proportion and APC model lag) that govern the breast cancer incidence can be different for black and white women. As such, determining the differences in percentage of aggressive tumors and the tumor growth rate could be especially useful in explaining the racial disparities in breast cancer mortality (see Fryback et al. [2006] for more details about UWBCS natural history parameters).

The calibration process of the race-specific UWBCS requires evaluation of 378,000 parameter combinations for each race. Therefore, it is crucial to use accelerated approaches for the calibration [Kong et al., 2009]. We refer to the calibrated versions of the UWBCS that replicates the observed outcomes for black and white women as UWBCS-black and UWBCS-white, respectively.

Parameter Name	Best fit in the original combined race UWBCS	Sampled parameter values for the calibration of race-specific model			Number of Parameter Values
		LB	UB	Increments	
Fraction of LMP tumors	0.42	0.00	0.60	0.10	7
In situ tumor boundary	0.95	0.95	0.95	-	1
Max LMP size	2.00	2.00	2.00	-	1
LMP dwell time	0.50	1.00	3.00	1.00	3
Onset proportion	0.90	0.80	1.00	0.05	5
APC model lag	3.00	1.00	7.00	2.00	4
Percentage of aggressive tumors	0.01	0.00	0.20	0.05	5
Percentage of highly aggressive tumors	0.02	0.00	0.10	0.025	5
Mean tumor growth	0.12	0.00	0.05	0.01	6
Variance tumor growth	0.01	0.00	0.05	0.01	6

Table 5.1: UWBCS natural history parameters.

5.2.5 Using active learning to speed-up the calibration process

The evaluation of a large number of parameter combinations can be avoided by identifying smaller subsets of the combinations that are more likely to produce desired outputs. Often evaluating only a small subset of input parameter combinations is sufficient for calibration, and the determination of such subsets is of key importance for the identification of natural history parameters. To this end, statistical models such as linear regression or machine learning methods such as artificial neural networks (ANNs) can be used as prediction models to determine promising subsets from the parameter space. For our problem, we use ANNs to expedite the calibration process. We develop an active learning algorithm that guides the ANN model to choose which parameter combinations to evaluate during the calibration.

5.2.5.1 Using a prediction model for calibration

We summarize the usage of a prediction model in the calibration process in five steps (see Figure 5.2). We start with identifying the set of parameter combinations from the parameter

space of unobservable parameters. Then, we randomly select a small subset of these possible parameter combinations and evaluate them via simulation to obtain a numerical score, which indicates how closely the output of each parameter combination matches target outcomes. We note that the choice of scoring approach is problem-specific and a generic prediction model performs independently of the scoring scheme. By using this small set of evaluated parameter combinations as our training set, we build a machine learning prediction model, such as an ANN, that predicts the numerical score associated with a given parameter combination. Then, we use this prediction model to find those combinations that have acceptable predicted scores. We refer to a score associated with a parameter combination as an acceptable score only if it falls into a predetermined range for the scores. Those parameter combinations with acceptable predicted scores are more likely to produce desired outputs from the simulation. Finally, we evaluate the parameter combinations with acceptable predicted scores using the simulation and identify the parameter combinations that best correspond to the values of the unobservable parameters.

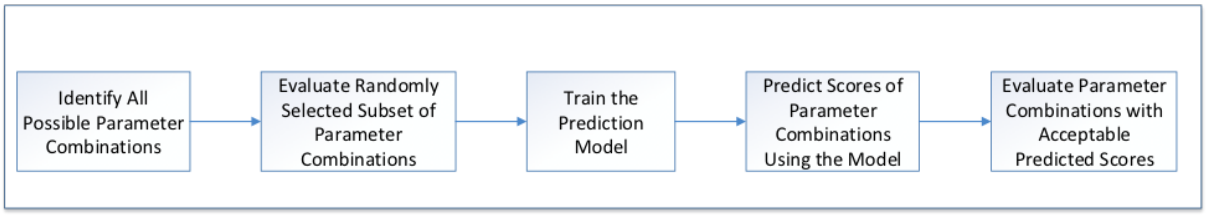

Figure 5.2: Overview of the use of a prediction model in simulation calibration.

5.2.5.2 Scores for parameter combinations

Scoring the parameter combinations is an important step of the calibration process. The score of a parameter combination is a measure of the difference between the simulation result and target outcome. A scoring scheme is typically model-specific and is affected by how modelers value different targets. A general method to determine the score of a parameter combination is to assign weights to each target and multiply those weights by the absolute or mean squared

error between the simulation result and the target outcome. For the simulation models that aim to capture the trends in the observed outcome (e.g., cancer incidence by calendar year or age), an alternative method is to identify acceptance envelopes around these outcomes and determine the scores using these envelopes. Whenever the simulation result corresponding to a target year or age falls outside of the envelope, it is a violation of the closeness measure and the score of the parameter combination is increased by one. Therefore, smaller scores represent better parameter combinations. Note that acceptance envelopes are usually formed by determining error margins around the target outcome and can be modified to better capture the trends in a given year or age group.

In order to assess the closeness of the output incidence generated by a given natural history parameter combination to the observed breast cancer incidence in the U.S., we generate envelopes around the observed incidence, which provide an acceptable error margin. Note that the error margins are determined according to expert opinion and they allow focusing on specific calendar years to capture the trends better for those years. We then count the number of times the output incidence falls outside of these envelopes for each breast cancer stage to assign a score representing the goodness of the output incidence. Because there are 4 incidence values for each cancer stage (in situ, localized, regional, distant), mortality values, and 36 years between 1975-2010, the worst score for an input parameter is 180. As the smaller scores indicate a better fit to the observed data, the best score is 0, which is obtained when all output incidence values are within the envelopes. We assume that a score less than 20 is an acceptable match.

5.2.5.3 ANNs as a prediction model

One way to build a prediction model is to use a linear regression model, which has the following functional form:

$$f(x) = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \cdots + \beta_n X_n, \tag{5.1}$$

where $X_1, \ldots, X_n$ denote n predictor variables (e.g., calibrated natural history parameters of the UWBCS), $f(X)$ denotes the estimated score of a given parameter combination, β_0 is a constant and $\beta_1, \ldots, \beta_n$ are the regression coefficients of the predictor variables $X_1, \ldots, X_n$. While linear regression is an easy-to-use prediction method, it will not be able to represent the underlying function accurately if the latter is nonlinear. Therefore, we consider artificial neural networks as a more expressive prediction model that is able to represent nonlinear relationships between the predictor variables. Note that there are also several other nonlinear regression models that consider the interactions between the predictor variables and we use the simple linear regression solely as a benchmark for the performance of other approaches presented in our study.

Artificial neural networks are information processing structures inspired by the function of biological neural networks. They consist of a set of interconnected units where each unit has a transfer function that computes and outputs a function of the values provided by the units that connect to it [Russell and Norvig, 2003]. Instead of the linear function used in the linear regression for the estimation of scores, ANNs also employ nonlinear functions as transfer functions at some units. The most commonly used transfer functions are sigmoid (logistic) functions, which can be given as $g(x) = \frac{1}{1+e^{-x}}$. Alternatively, hyperbolic tangent functions or Gaussian functions can be used as transfer functions. Note that single layer ANNs with a linear transfer function $f(X)$ given in equation 5.1 are equivalent to linear regression models. A graphical representation of a three layer neural network with two inputs and one output is shown in Figure 5.3.

ANNs vary in the way neurons are connected and inputs are processed. In this study, we use a 3-layer feed-forward ANN trained using a backpropagation learning algorithm with a learning rate 0.3 and momentum 0.2 [Russell and Norvig, 2003]. The layers include an input layer representing the natural history parameters, a hidden layer with 21 hidden units, and an output layer with a single unit generating an estimate for the score. Our ANN model aims to minimize the sum of squared errors between predicted scores and actual scores and it uses the

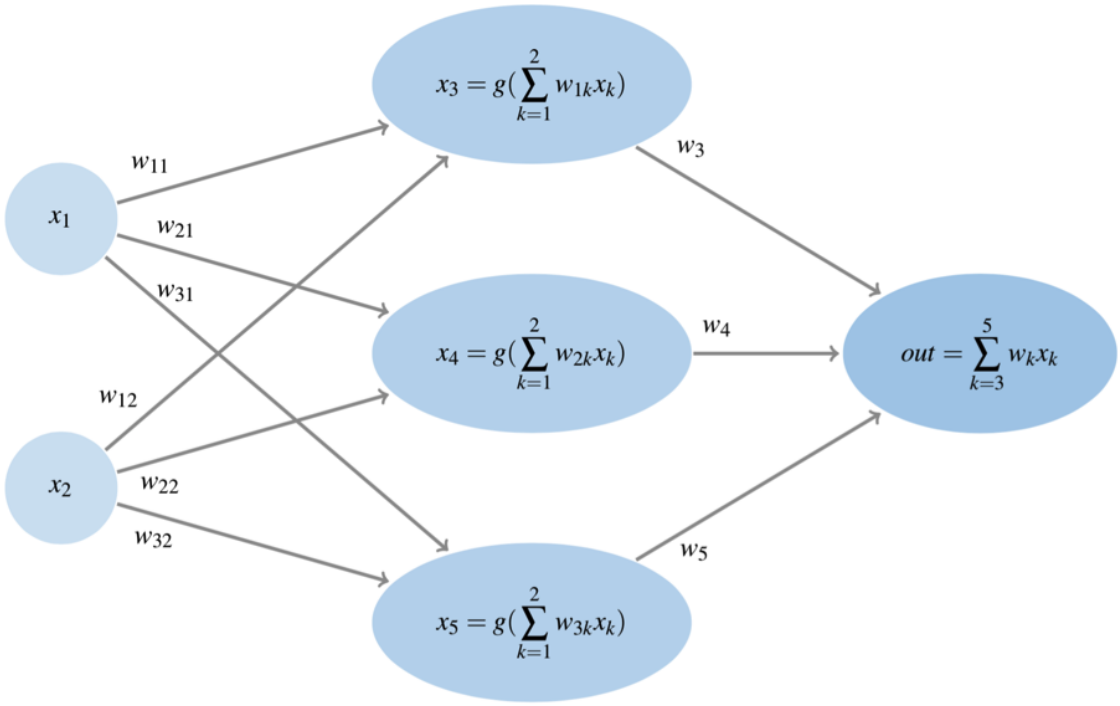

Figure 5.3: An illustrative graph showing a feed-forward ANN with an input layer with 2 inputs, a hidden layer with 3 hidden units, and one output unit.

sigmoid function as the transfer function in the hidden units and a linear transfer function in the output unit. Note that there are no exact rules in determining an ANN model's parameters. For instance, learning rate and momentum parameters controls the speed of convergence of the backpropagation algorithm and both parameters can take any value between 0 and 1. Choosing higher values for these parameters may lead to divergence of the algorithm, whereas smaller values usually cause a slow convergence. We obtain the final values for such parameters by experimenting with the most commonly used values in the literature.

We also use ensemble methods, more specifically bagging [Breiman, 1996], to improve the predictive accuracy of the ANN model (see Appendix C.1 for more details about ensemble

methods). We refer to the ensemble of ANN models constructed using bagging as bagANN. Also, in keeping with the machine learning terminology, we refer a parameter combination as an instance, and individual input parameters as features. We refer to a parameter combination with an unknown score as an unscored instance.

5.2.5.4 Active learning for calibration

Supervised learning methods such as ANNs use randomly selected instances as training sets to build a prediction model. However, it is often possible to achieve a higher level of predictive accuracy by having the learning method actively sample instances to be included in the training set. The process that interactively samples the instances is known as active learning [Settles, 2012]. The difficulty of obtaining scores via simulation for every parameter combination makes active learning attractive for the calibration problem. Therefore, to increase the prediction power of the bagANN model and to decrease the total number of simulation runs required, we develop an active learning algorithm that iteratively selects promising parameter combinations from the set of all combinations and evaluates only these parameter combinations using the simulation model.

Figure 5.4 summarizes the active learning approach that is used in this study. First, we train the bagANN model using a small number of evaluated parameter combinations. Then, we use this bagANN model to select a batch of instances with low predicted scores from a large pool of unscored instances and evaluate these instances via the simulation model. If the stopping criterion is not met, we retrain the bagANN model with the selected instances and repeat the process. We use a clustering approach to select the batch of instances with low predicted scores to ensure that the putatively low-scoring instances in the batch represent the different regions of the parameter space. For clustering the instances, we use hierarchical clustering with Euclidean distance and single linkage measures. Moreover, we augment the batch that is used to update the bagANN model with a set of randomly selected parameter combinations. This step further

assures that the parameter space is broadly sampled to train the bagANN model. We present a more detailed version of our active learning algorithm in the appendix (see Appendix C.2).

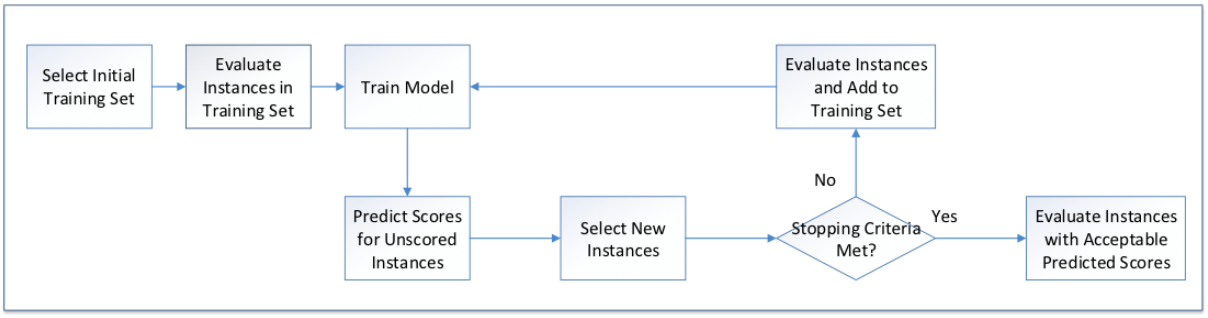

Figure 5.4: Flowchart of the simulation model calibration process using active learning.

An important element of the active learning algorithm is the stopping criterion [Vlachos, 2008]. For our method, we use a stopping criterion that detects when the active learning process is providing diminishing returns in identifying acceptable parameter combinations. We first determine a warm-up period for the active learning algorithm, which is the minimum number of iterations required to stop sampling. Then, we stop the active learning algorithm when no acceptable instance is added to the training set in k consecutive iterations, where k is a user-defined parameter.

5.2.5.5 Estimating the number of simulation runs

After training the bagANN model on a randomly selected subset of parameter combinations, we can test the accuracy of the bagANN model on an independently selected test set. This predictive accuracy information serves as a performance measure for our prediction model. In addition, the test set can be used to estimate the total number of simulation runs needed to obtain a reasonably high portion of all good parameter combinations. Algorithm 5.1 summarizes this simple procedure to estimate the average number of simulation runs to perform using the prediction model.

Algorithm 5.1: Calculate Expected Number of Runs

Input: test instances T, all instances S, threshold score δ

1. $T' \leftarrow \{i \in T \ : \ \ell(i) \leq \delta\}$
2. $\Delta \leftarrow \max_{i \in T'} \{f(i)\}$
3. $ENRuns \leftarrow \sum_{i \in S} I_{\{f(i) \leq \delta\}}$

In this algorithm, $\ell(i)$ represents the score of each instance in set T, $f(i)$ represents the predicted value for instance $i \in T$, and $I_{\{.\}}$ is the indicator function. Here, we first form a subset of test instances T' from the instances in the test set T that have scores lower than the predetermined threshold score δ. Then, we find the maximum predicted score Δ for those instances in the set T'. Finally, we count the number of instances with a predicted score lower than Δ in the set S, which is the set consisting of all instances.

5.3 Results

5.3.1 Performance of prediction models for the calibration

We only use the calibration process for the UWBCS-black to test the performance of the prediction models, as the two calibration processes are very similar. We assume that all of the features (individual parameters of a parameter combination) have nominal values in our experiments. We observe that changing the nominal features to numeric features does not lead to more accurate prediction models, which is also observed in several other studies [Maslove et al., 2013].

In our computational experiments, we first demonstrate the performance of the prediction models in predicting the scores of parameter combinations and then examine the additional benefits of using a prediction model in the active learning setting. Therefore, we first compare the bagANN model with a simple linear regression model and a single ANN model. For this experiment, we start with evaluating 30,000 parameter combinations using the simulation model

and use those as our training and test sets according to a 60/40 ratio (i.e. 60% of 30,000 are used for training and 40% of 30,000 are used for testing). As comparison measures, we use mean absolute error and root mean squared error values of the prediction models. We also report the predictive accuracy of the trained models on the instances with low actual scores by only considering the instances in the test set that have actual scores lower than some predetermined value. We use 30 as a threshold score to determine such instances and refer to these instances as low-scored instances.

In our experiments with the active learning algorithm, we start by evaluating 1000 parameter combinations and include these instances in a set L. At each iteration of the algorithm, we add 200 new instances to L. Half of those instances are selected based on predicted scores and hierarchical clustering, and the rest were randomly selected to minimize bias in the prediction model. In order to obtain an estimate of the number of acceptable parameter combinations identified after each iteration of the active learning algorithm, we count the number of acceptable parameter combinations that are in set L and among the set of 1000 instances selected from unscored instances based on their predicted scores. As a stopping criterion in the active learning experiments, we either use a bound on the number of instances in L or use 5 iterations for the *warm-up* period and $k = 3$ for the number of consecutive iterations without new acceptable instances added to the training set.

We carry out our experiments with the prediction models using the WEKA (version 3.6.5) JAVA library. Experiments are run on Intel Xeon 2.27 GHz processor with 12GB RAM running Windows 7.

5.3.1.1 Comparison of the prediction models

Table 1 reports the performance of the prediction models in a standard supervised learning setting where we use a single training set to train a model, and then make predictions based on this model. We also report the estimates on the additional number of evaluations required for

various prediction models, which are obtained using Algorithm 1. Our results show that the ANNs have significantly better predictive accuracy than linear regression (LinR) as observed in mean absolute error (MAE) and root mean squared error (RMSE) values as well as the number of additional runs required by each model. Besides, improvements in the predictive accuracy are more apparent in MAE and RMSE values obtained for the low-scored instances. We recognize that the simple linear regression model does not account for the interactions between different features. However, ANNs use nonlinear activation functions to overcome this limitation, which also helps ANNs to achieve better predictive accuracy for our problem.

Using bagging to construct ensembles of ANNs further improves predictive accuracy especially in terms of the additional number of runs required by each model. While the ANN model alone suggested that an additional 4626 simulation evaluations would generate all the acceptable parameter combinations, bagANN model achieves the same level of performance through the evaluation of only 2741 additional parameter combinations. An experiment regarding how bagANN performance varies as a function of the training set size is available in Appendix C.3.

Training Method	MAE	RMSE	MAE-low	RMSE-low	ENRuns ($\delta = 15$)	Time (in sec)
LinR	8.16	10.97	33.45	34.08	78754(99%)	1
ANN	3.26	4.36	7.00	8.42	4626(100%)	290
bagANN	2.26	2.97	5.41	6.46	2741(100%)	2547

Table 5.2: Comparison of the different prediction models for the UWBCS calibration problem.

5.3.1.2 Benefits of using active learning

Our next experiment tests the benefit of using an active learning framework to select the instances that are provided to the bagANN for training. In this experiment, we compare our active learning algorithm to a method that randomly selects the instances that are added to the training set. For both cases, we repeat the algorithm for 30 different initial training sets each including 2000 parameter combinations and stop the algorithm when our training set reaches to 5000

instances. Figure 5.5 shows that the active learning algorithm can locate all acceptable parameter combinations by evaluating only a fraction of all instances whereas the random selection method needs to evaluate significantly more instances to achieve similar performance. In fact, random selection method does not achieve the same level of performance with the active learning algorithm even if we continue sampling instances until the training set reaches 10,000 instances. Error bars in Figure 5.5 indicate that active learning results are slightly more variable with respect to the initial training set compared to random selection. We provide one-way sensitivity analysis for the several model parameters in the appendix (see Appendix C.4).

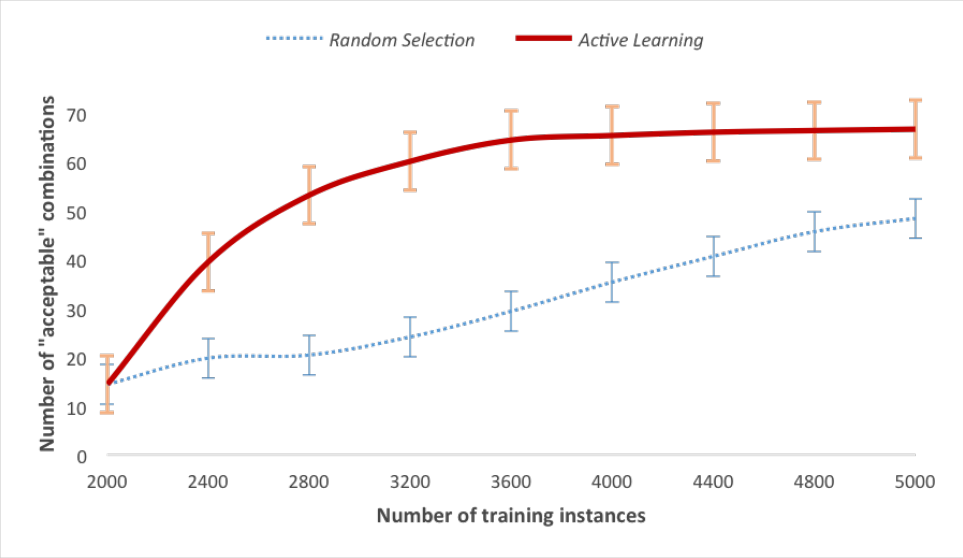

Figure 5.5: Learning curves showing the number of acceptable parameter combinations found as a function of the number of training instances: active learning compared with random selection.

5.3.2 Model fit to the U.S.-observed breast cancer incidence and mortality rates

We use the calibrated natural history parameters to obtain race-specific UWBCS, namely UWBCS-black and UWBCS-white. Figure 5.6a shows the comparison of UWBCS-black with the U.S.-observed incidence and mortality rates. While UWBCS-black successfully replicates the breast

cancer incidence for each cancer stage, the fit for the U.S.-observed mortality rates is not as good. However, the simulation model successfully captures the general trends in breast cancer mortality for black women. UWBCS-white performs better compared to UWBCS-white in terms of fit to the U.S.-observed rates (see Figure 5.6b).

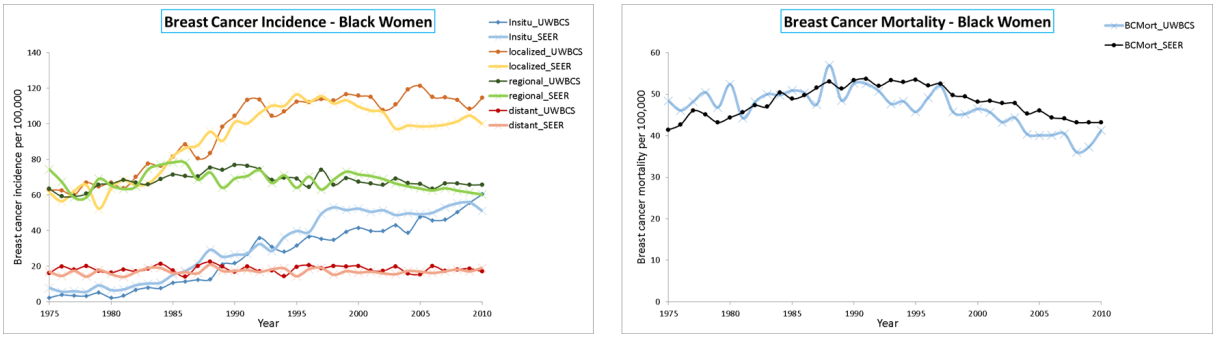

(a) Model fit for black women

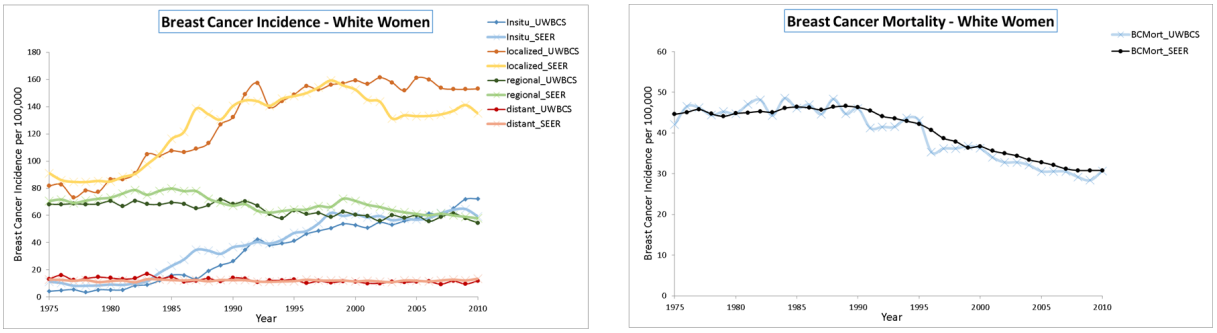

(b) Model fit for white women

Figure 5.6: Comparison of UWBCS results with U.S.-observed breast cancer incidence and mortality rates

5.3.3 Factors leading to higher mortality rates in black women

Table 5.3 compares age-adjusted breast cancer mortality in 1975 to 2010 among white women and black women to predictions from series of models with UWBCS-white parameter values sequentially replaced by black values. Note that we also provide the results from two other simulation models (MISCAN and SPECTRUM) that have been used byvan Ravesteyn et al. [2011]

and compare their findings with UWBCS. Of the difference between observed mortality and predicted mortality after taking into account lower incidence between among blacks, natural history explained 65.9%, screening use 3.9%, and use of adjuvant treatment therapy 20.7%, leaving 9% unexplained in UWBCS. In explaining the mortality difference, UWBCS attributes more weight to natural history parameters compared to MISCAN and SPECTRUM. However, van Ravesteyn et al. [2011] note that MISCAN and SPECTRUM fail to replicate observed breast cancer mortality rates for black women and significant proportion of differences cannot be explained by their models (38-46%). As a result, UWBCS does a better job in quantifying the impact of each factor on explaining the higher mortality rates in black women.

	natural history	natural history and **screening**	natural history, screening, and **treatment**	**all** (UWBCS-black)	**unexplained**
MISCAN	26.0%	8.0%	19.0%	54.0%	46.0%
SPECTRUM	44.0%	7.0%	11.0%	62.0%	38.0%
UWBCS	65.9%	3.9%	20.7%	91.0%	9.0%

Table 5.3: Effect of sequential replacement of parameters for black women in the UWBCS-white on the breast cancer mortality rate for black women.

5.4 Concluding Remarks

We use UWBCS, a previously validated discrete-event stochastic simulation model of breast cancer, to evaluate the differences in breast cancer incidence and mortality between black and white women. Model parameters are estimated for the two racial groups while accounting for differences in mammography uptake and treatment dissemination, and the resulting models, UWBCS-black and UWBCS-white, closely replicates the observed breast cancer outcomes by SEER [SEER, 2008a,b]. We observe that breast cancer natural history parameters are a major driver for race-specific differences in mortality, which is also noted by Batina et al. [2013] and van

Ravesteyn et al. [2011]. Therefore, targeted prevention and detection strategies that go beyond equalizing access to mammography may be needed to eliminate racial disparities. Furthermore, our results show that diminished use of adjuvant treatment is much more impactful on increased breast cancer mortality among black women compared to inadequate mammography utilization.

We also provide a novel framework to speed up the simulation calibration. In this study, we provide a method to determine the most promising input values so that calibration can be completed by evaluating only a small subset of all parameter combinations. For this purpose, we use several machine learning approaches and provide several key statistics on the predictive accuracy of each method. Our results show that the computational burden of the calibration may be reduced significantly and the calibration is feasible even for complex simulation models such as UWBCS.

To the best of our knowledge, the study by Kong et al. [2009] is the only calibration work that described a framework to reduce the number of parameter combinations to be evaluated. More specifically, Kong et al. [2009] use metaheuristics such as simulated annealing and genetic algorithms to search the parameter space to find the best parameter combination for a lung cancer simulation model. They also provide a method for measuring the closeness of the actual calibration targets and simulation outcomes. While their approach effectively avoids evaluating all parameter combinations, it is able to identify only one (or a small number of) acceptable parameter combination that matches the target outcome without actually assessing the quality of the outputs for other parameter combinations. On the other hand, identifying all or a big portion of acceptable parameter combinations can help guard against non-identifiability problems in calibration. As our approach provides estimations on outcomes of any parameter combination, one can take advantage of all parameter combinations to better understand the effects of parameter uncertainty on model conclusions.

Our active learning algorithm and bagANN prediction model do not guarantee finding every acceptable parameter combination and we are able to report the percentage of acceptable

parameter combinations found by our method only because we have actually evaluated all the parameter combinations beforehand. In addition, because the active learning algorithm stops when there are no new acceptable parameter combinations found in κ consecutive iterations, the number of acceptable parameter combinations obtained by the active learning algorithm can be used as an estimate on the total number of acceptable parameter combinations in the parameter space. Note that there are no exact procedures for the selection of the warm-up period and the κ value and these parameters may significantly affect the performance of the prediction model. For instance, if we increase the length of warm-up period (or κ), we expect our bagANN model to predict scores more accurately. However, increased values for these parameters also lead to more simulation runs. Therefore, it is important to find the right balance between the prediction model's predictive accuracy and the total number of simulation runs required to obtain this accuracy. We perform a one-way sensitivity analysis to demonstrate the impact of stopping criteria parameters on the performance of active learning algorithm.

While we demonstrate our new calibration method using only UWBCS, it can easily be applied to other simulation models that require the evaluation of a large number of parameter combinations for the calibration. In general, calibration processes of the simulation models require going through similar steps. After deciding on model-specific scoring scheme and determining the set of parameter combinations, modelers can use one of the two approaches presented in his study. The first approach is summarized in Figure 5.2 and involves constructing a prediction model using the initial simulation evaluations and estimating the scores of the remaining parameter combinations. While this approach is easier to implement, it is difficult to determine ideal training set size and threshold score δ, which is used to determine which parameter combinations to evaluate. The second approach is to use an active learning algorithm, which starts with a prediction model that is constructed using a very small training set and reconstructs this model by iteratively adding more instances to the training set as shown in Figure 5.4. As we start with a very small training set and only add those instances to the

training set that lead to a better prediction model, active learning approach generally requires significantly less simulation evaluations and also includes a natural stopping criterion. Modelers may choose between these two approaches according to the complexity of their model and available computational resources. In addition, because of the similarities in the calibration processes, we expect our approaches to perform well for other simulation models. However, model-specific features such as the number of unobservable parameters and scoring scheme may affect some parameters of our algorithms (e.g., initial training set size and stopping criteria parameters of the active learning algorithm) as well as the type of the machine learning algorithm used as the prediction model.

Several parameters related to prediction models and the active learning algorithm are determined based on expert opinions and empirical analysis. We conduct numerical experiments to show the effects of some of these parameters such as initial training set size and threshold score δ on model performance. On the other hand, we rely on the most commonly used methods in the literature and our preliminary analysis for the selection of many other parameters such as type of the activation function used in the ANN and type of the clustering approach employed in the active learning algorithm. In addition, we recognize that other prediction models, including different types of regression models and ensemble methods, may perform better than bagANN for a general simulation model.

BIBLIOGRAPHY

ACS. *Cancer Facts and Figures 2011*. American Cancer Society, Atlanta, GA, 2011. (Cited on pages 1 and 2.)

O. Alagoz, T. Ayer, and F. S. Erenay. Operations research models for cancer screening. In J. J. Cochran, L. A. Cox Jr., P. Keskinocak, J. P. Kharoufeh, and J. C. Smith, editors, *Wiley Encyclopedia of Operations Research and Management Science*. John Wiley & Sons, Hoboken, NJ, 2011. (Cited on page 12.)

K. Amend, D. Hicks, and C. B. Ambrosone. Breast cancer in african-american women: differences in tumor biology from european-american women. *Cancer Research*, 66(17):8327–8330, 2006. (Cited on page 8.)

E. Arias. United states life tables, 2010. *National Vital Statistics System*, 63(7):1–63, 2014. (Cited on page 62.)

T. Ayer, O. Alagoz, and N. K. Stout. A POMDP approach to personalize mammography screening decisions. *Operations Research*, 60(5):1019–1034, 2012. (Cited on pages vi, 3, 11, 13, and 121.)

T. Ayer, O. Alagoz, N. K. Stout, and E. S. Burnside. Heterogeneity in women's adherence and its role on optimal breast cancer screening policies. *Management Science*, 62:1339–1362, 2016. (Cited on page 13.)

M. U. Ayvaci, O. Alagoz, and E. S. Burnside. The effect of budgetary restrictions on breast cancer diagnostic decisions. *Manufacturing & Service Operations Management*, 14(4):600–617, 2012. (Cited on pages 2 and 13.)

S. BahadurSingh, R. Maharaj, P. Harnarayan, S. O. Cawich, M. Yearwood, and V. Naraynsingh. Mammographic screening: Is it relevant to developing countries? *Current Medicine Research and Practice*, 4(4):168–170, 2014. (Cited on page 6.)

M. Banerjee, J. George, C. Yee, W. Hryniuk, and K. Schwartz. Disentangling the effects of race on breast cancer treatment. *Cancer*, 110(10):2169–2177, 2007. (Cited on page 14.)

M. B. Barton, R. Harris, and S. W. Fletcher. Does this patient have breast cancer?: The screening clinical breast examination: Should it be done? How? *JAMA*, 282(13):1270–1280, 1999. (Cited on pages 32 and 61.)

N. G. Batina, A. Trentham-Dietz, R. E. Gangnon, B. L. Sprague, M. A. Rosenberg, N. K. Stout, D. G. Fryback, and O. Alagoz. Variation in tumor natural history contributes to racial disparities in breast cancer stage at diagnosis. *Breast Cancer Research and Treatment*, 138(2):519–528, 2013. (Cited on pages 15, 78, and 94.)

K. R. Bauer, M. Brown, R. D. Cress, C. A. Parise, and V. Caggiano. Descriptive analysis of estrogen receptor (er)-negative, progesterone receptor (pr)-negative, and her2-negative invasive breast cancer, the so-called triple-negative phenotype. *Cancer*, 109(9):1721–1728, 2007. (Cited on page 14.)

N. Baxter et al. Preventive health care, 2001 update: Should women be routinely taught breast self-examination to screen for breast cancer? *Canadian Medical Association Journal*, 164(13): 1837–1846, 2001. (Cited on pages 32 and 61.)

W. A. Berg. How well does supplemental screening magnetic resonance imaging work in high-risk women? *Journal of Clinical Oncology*, 32(21):2193–2196, 2014. (Cited on page 5.)

W. A. Berg, J. D. Blume, A. M. Adams, R. A. Jong, R. G. Barr, et al. Reasons women at elevated risk of breast cancer refuse breast MR Imaging screening: Acrin 6666 1. *Radiology*, 254(1):79–87, 2009. (Cited on page 33.)

W. A. Berg, Z. Zhang, D. Lehrer, R. A. Jong, E. D. Pisano, et al. Detection of breast cancer with addition of annual screening ultrasound or a single screening mri to mammography in women with elevated breast cancer risk. *JAMA*, 307(13):1394–1404, 2012. (Cited on pages 5 and 17.)

Berkeley. Berkeley Mortality Database for the United States, 2016. URL `http://www.demog.berkeley.edu/~bmd/states.html`. (Cited on page 80.)

N. Biller-Andorno and P. Jüni. Abolishing mammography screening programs? a view from the swiss medical board. *New England Journal of Medicine*, 370(21):1965–1967, 2014. (Cited on page 1.)

L. Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996. (Cited on pages 86 and 122.)

E. S. Burnside, J. Belkora, and L. Esserman. The impact of alternative practices on the cost and quality of mammographic screening in the united states. *Clinical Breast Cancer*, 2(2):145–152, 2001. (Cited on page 3.)

L. A. Carey, C. M. Perou, C. A. Livasy, L. G. Dressler, D. Cowan, et al. Race, breast cancer subtypes, and survival in the carolina breast cancer study. *JAMA*, 295(21):2492–2502, 2006. (Cited on page 14.)

A. Cassandra, M. L. Littman, and N. L. Zhang. Incremental pruning: A simple, fast, exact method for partially observable Markov decision processes. In *Proceedings of the Thirteenth conference on Uncertainty in artificial intelligence*, pages 54–61. Morgan Kaufmann Publishers Inc., 1997. (Cited on page 12.)

CDC. Centers for Disease Control and Prevention, 2016. URL `http://www.cdc.gov/nchs/`. (Cited on page 80.)

Census. United States Census Bureau Population Estimates, 2016. URL `http://www.census.gov/popest/`. (Cited on page 80.)

S. Chen and G. Parmigiani. Meta-analysis of brca1 and brca2 penetrance. *Journal of Clinical Oncology*, 25(11):1329–1333, 2007. (Cited on page 34.)

H. Cheng. *Algorithms for Partially Observable Markov Decision Processes*. PhD thesis, University of British Columbia, 1988. (Cited on page 12.)

J. Chhatwal, O. Alagoz, and E. S. Burnside. Optimal breast biopsy decision-making based on mammographic features and demographic factors. *Operations Research*, 58(6):1577–1591, 2010. (Cited on pages 2, 13, and 34.)

China Statistical Yearbook. National Bureau of Statistics of China, 2015. Available from http://www.stats.gov.cn/tjsj/ndsj/2014/indexeh.htm. Accessed on February 26, 2015. (Cited on page 6.)

R. T. Chlebowski, Z. Chen, G. L. Anderson, T. Rohan, A. Aragaki, et al. Ethnicity and breast cancer: factors influencing differences in incidence and outcome. *JNCI*, 97(6):439–448, 2005. (Cited on pages 8 and 14.)

K. A. Cronin, D. L. Miglioretti, M. Krapcho, B. Yu, B. M. Geller, P. A. Carney, T. Onega, E. J. Feuer, N. Breen, and R. Ballard-Barbash. Bias associated with self-report of prior screening mammography. *Cancer Epidemiology Biomarkers & Prevention*, 18(6):1699–1705, 2009. (Cited on page 80.)

T. G. Dietterich. Ensemble methods in machine learning. In *Multiple classifier systems*, pages 1–15. Springer, 2000. (Cited on page 122.)

M. F. Drummond. *Methods for the economic evaluation of health care programmes*. Oxford university press, 2005. (Cited on page 72.)

J. N. Eagle. The optimal search for a moving target when the search path is constrained. *Operations Research*, 32(5):1107–1115, 1984. (Cited on page 11.)

C. C. Earle, R. H. Chapman, C. S. Baker, C. M. Bell, P. W. Stone, E. A. Sandberg, and P. J. Neumann. Systematic overview of cost-utility assessments in oncology. *Journal of Clinical Oncology*, 18 (18):3302–3317, 2000. (Cited on page 62.)

James E Eckles. Optimum maintenance with incomplete information. *Operations Research*, 16(5): 1058–1067, 1968. (Cited on page 11.)

H. Ellis, M. Jiang, and R. B. Corotis. Inspection, maintenance, and repair with partial observability. *Journal of Infrastructure Systems*, 1(2):92–99, 1995. (Cited on page 11.)

J. G. Elmore, M. B. Barton, V. M. Moceri, S. Polk, Arena P. J., and Fletcher S. W. Ten-year risk of false positive screening mammograms and clinical breast examinations. *New England Journal of Medicine*, 338:1089–1096, 1998. (Cited on page 1.)

J. G. Elmore, C. Y. Nakano, H. M. Linden, L. M. Reisch, J. Z. Ayanian, and E. B. Larson. Racial inequities in the timing of breast cancer detection, diagnosis, and initiation of treatment. *Medical Care*, 43(2):141–148, 2005. (Cited on pages 14, 32, and 61.)

F. S. Erenay. *Optimal Screening Policies for Colorectal Cancer Prevention and Surveillance*. PhD thesis, UW-Madison, 2010. (Cited on page 51.)

F. S. Erenay, O. Alagoz, and A. Said. Optimizing colonoscopy screening for colorectal cancer prevention and surveillance. *Manufacturing & Service Operations Management*, 16(3):381–400, 2014. (Cited on page 13.)

S. A. Fedewa, E. M. Ward, A. K. Stewart, and S. B. Edge. Delays in adjuvant chemotherapy treatment among patients with breast cancer are more likely in african american and hispanic populations: A national cohort study 2004-2006. *Journal of Clinical Oncology*, pages JCO–2009, 2010. (Cited on page 14.)

D. G. Fryback, N. K. Stout, M. A. Rosenberg, A. Trentham-Dietz, V. Kuruchittham, and P. L. Remington. The Wisconsin breast cancer epidemiology simulation model. *JNCI Monographs*, 2006(36):37–47, 2006. (Cited on pages 31, 61, 79, 81, and 121.)

M. H. Gail and B. Rimmer. Risk-based recommendations for mammographic screening for women in their forties. *Journal of Clinical Oncology*, 169:3105–3114, 1998. (Cited on page 3.)

M. H. Gail, J. P. Costantino, J. Bryant, R. Croyle, L. Freedman, K. Helzlsouer, and V. Vogel. Weighing the risks and benefits of tamoxifen treatment for preventing breast cancer. *JNCI*, 91 (21):1829–1846, 1999. (Cited on page 61.)

A. Ghafoor, A. Jemal, V. Cokkinides, C. Cardinez, T. Murray, A. Samuels, and M. J. Thun. Cancer statistics for african americans. *CA: A Cancer Journal for Clinicians*, 52(6):326–341, 2002. (Cited on page 7.)

O. Gilbar. Coping with threat: Implications for women with a family history of breast cancer. *Psychomatics*, 39:229–329, 1998. (Cited on pages 3 and 7.)

P. M. Granitto, P. F. Verdes, and H. A. Ceccatto. Neural network ensembles: evaluation of aggregation algorithms. *Artificial Intelligence*, 163(2):139–162, 2005. (Cited on page 122.)

Barry I Graubard, Andrew N Freedman, and Mitchell H Gail. Five-year and lifetime risk of breast cancer among us subpopulations: implications for magnetic resonance imaging screening. *Cancer Epidemiology Biomarkers & Prevention*, 19(10):2430–2436, 2010. (Cited on page 62.)

A. M. Gray, P. M. Clarke, J. L. Wolstenholme, and S. Wordsworth. *Applied methods of cost-effectiveness analysis in healthcare*. Oxford University Press, 2010. (Cited on page 47.)

E. D. Güneş, E. L. Örmeci, and D. Kunduzcu. Preventing and diagnosing colorectal cancer with a limited colonoscopy resource. *Production and Operations Management*, 24(1):1–20, 2015. (Cited on page 14.)

M. Hauskrecht and H. Fraser. Planning treatment of ischemic heart disease with partially observable Markov decision processes. *Artificial Intelligence in Medicine*, 18(3):221–244, 2000. (Cited on page 11.)

D. Hershman, R. McBride, J. S. Jacobson, L. Lamerato, K. Roberts, V. R. Grann, and A. I. Neugut. Racial disparities in treatment and survival among women with early-stage breast cancer. *Journal of Clinical Oncology*, 23(27):6639–6646, 2005. (Cited on page 15.)

D. L. Hershman, J. M. Unger, W. E. Barlow, L. F. Hutchins, S. Martino, C. K. Osborne, R. B. Livingston, and K. S. Albain. Treatment quality and outcomes of african american versus white breast cancer patients: retrospective analysis of southwest oncology studies s8814/s8897. *Journal of Clinical Oncology*, 27(13):2157–2162, 2009. (Cited on page 15.)

T. R. Holford, K. A. Cronin, A. B. Mariotto, and E. J. Feuer. Changing patterns in breast cancer incidence trends. *Monographs - National Cancer Institue*, 36:19, 2006. (Cited on page 80.)

N. Howlader, A. M. Noone, M. Krapcho, N. Neyman, R. Aminou, et al. Seer cancer statistics review, 1975–2008. *Bethesda, MD: National Cancer Institute*, 19, 2011. (Cited on page 7.)

C. Hu, W. S. Lovejoy, and S. L. Shafer. Comparison of some suboptimal control policies in medical drug therapy. *Operations Research*, 44(5):696–709, 1996. (Cited on page 11.)

IRQN. IRQN welcomes the Chinese Society of Radiology, 2015. Available from http://irqn.org/media/news.htm. Accessed on February 26, 2015. (Cited on page 6.)

J. S. Ivy. Can we do better? Optimization models for breast cancer screening. In *Handbook of Optimization in Medicine*, pages 1–28. Springer, 2009. (Cited on page 12.)

U. W. Jayasinghe, R. Taylor, and J. Boyages. Is age at diagnosis an independent prognostic factor for survival following breast cancer? *ANZ Journal of Surgery*, 75:762–767, 2005. (Cited on page 2.)

A. R. Jazieh and C. R. Buncher. Racial and age-related disparities in obtaining screening mammography: results of a statewide database. *Southern Medical Journal*, 95(10):1145–1149, 2002. (Cited on page 14.)

S. A. Joslyn. Racial differences in treatment and survival from early-stage breast carcinoma. *Cancer*, 95(8):1759–1766, 2002. (Cited on page 15.)

L. Kallenberg and M. Cornelis. Survey of linear programming for standard and nonstandard Markovian control problems. Part I: Theory. *Zeitschrift für Operations Research*, 40(1):1–42, 1994. (Cited on page 55.)

K. Kerlikowske, D. Grady, S. M. Rubin, C. Sandrock, and V. L. Ernster. Efficacy of screening mammography: A meta-analysis. *JAMA*, 273:149–154, 1995. (Cited on page 1.)

K. Kerlikowske, P. A. Carney, and B. Geller. Performance of screening mammography among women with and without a first-degree relative with breast cancer. *Annals of Internal Medicine*, 133:855–863, 2000. (Cited on pages 2 and 61.)

K. Kerlikowske, R. A. Hubbard, D. L. Miglioretti, B. M. Geller, B. C. Yankaskas, et al. Comparative effectiveness of digital versus film-screen mammography in community practice in the united states: a cohort study. *Annals of Internal Medicine*, 155(8):493–502, 2011. (Cited on page 4.)

C. N. Klabunde, R. Ballard-Barbash, and International Breast Cancer Screening Network. Evaluating population-based screening mammography programs internationally. In *Seminars in Breast Disease*, volume 10, pages 102–107. Elsevier, 2007. (Cited on page 2.)

F. M. Knuttel, G. L. G. Menezes, M. A. A. J. van den Bosch, K. G. A. Gilhuijs, and N. H. G. M. Peters. Current clinical indications for magnetic resonance imaging of the breast. *Journal of Surgical Oncology*, 110(1):26–31, 2014. (Cited on page 32.)

C. Y. Kong, P. M. McMahon, and G. S. Gazelle. Calibration of disease simulation model using an engineering approach. *Value in Health*, 12(4):521–529, 2009. (Cited on pages 81 and 95.)

C. H. Lee, D. D. Dershaw, D. Kopans, P. Evans, B. Monsees, et al. Breast cancer screening with imaging: Recommendations from the Society of Breast Imaging and the ACR on the use of mammography, breast MRI, breast ultrasound, and other technologies for the detection of clinically occult breast cancer. *Journal of the American College of Radiology*, 7(1):18–27, 2010. (Cited on page 2.)

C. I. Lee, M. Cevik, O. Alagoz, B. L. Sprague, A. N. A. Tosteson, et al. Comparative effectiveness of combined digital mammography and tomosynthesis screening for women with dense breasts. *Radiology*, 2014. (Cited on pages 4, 5, and 14.)

C. P. Lee, G. M. Chertow, and S. A. Zenios. An empiric estimate of the value of life: Updating the renal dialysis cost-effectiveness standard. *Value in Health*, 12(1):80–87, 2009. (Cited on page 51.)

C. I. Li, K. E. Malone, and J. R. Daling. Differences in breast cancer stage, treatment, and survival by race and ethnicity. *Archives of Internal Medicine*, 163(1):49–56, 2003. (Cited on page 15.)

Y. Li, M. Zhu, R. Klein, and N. Kong. Using a partially observable markov chain model to assess colonoscopy screening strategies–a cohort study. *European Journal of Operational Research*, 238 (1):313–326, 2014. (Cited on page 13.)

W. S. Lovejoy. A survey of algorithmic methods for partially observed Markov decision processes. *Annals of Operations Research*, 28:47–66, 1991a. (Cited on pages 12 and 25.)

W. S. Lovejoy. Computationally feasible bounds for partially observed Markov decision processes. *Operations Research*, 39(1):162–175, 1991b. (Cited on pages 12, 25, 52, 53, 111, 112, and 113.)

M. Madadi, S. Zhang, and L. M. Henderson. Evaluation of breast cancer mammography screening policies considering adherence behavior. *European Journal of Operational Research*, 247(2):630–640, 2015. (Cited on page 13.)

L. M. Maillart, J. S. Ivy, K. Diehl, and S. Ransom. Assessing dynamic breast cancer screening policies. *Operations Research*, 56:1411–1427, 2008. (Cited on page 13.)

Lisa M Maillart. Maintenance policies for systems with condition monitoring and obvious failures. *IIE Transactions*, 38(6):463–475, 2006. (Cited on page 11.)

J. S. Mandelblatt, M. E. Wheat, M. Monane, R. D. Moshief, J. P. Hollenberg, and J. Tang. Breast cancer screening for elderly women with and without comorbid conditions: A decision analysis model. *Annals of Internal Medicine*, 116(9):722–730, 1992. (Cited on pages 34 and 62.)

J. S. Mandelblatt, N. K. Stout, C. B. Schechter, J. J. van den Broek, D. L. Miglioretti, et al. Collaborative modeling of the benefits and harms associated with different us breast cancer screening strategies. *Annals of Internal Medicine*, 2016. (Cited on pages 65 and 121.)

A. Mariotto, E. J. Feuer, L. C. Harlan, L. Wun, K. A. Johnson, and J. Abrams. Trends in use of adjuvant multi-agent chemotherapy and tamoxifen for breast cancer in the united states: 1975–1999. *JNCI*, 94(21):1626–1634, 2002. (Cited on page 80.)

A. B. Mariotto, E. J. Feuer, L. C. Harlan, and J. Abrams. Dissemination of adjuvant multiagent chemotherapy and tamoxifen for breast cancer in the united states using estrogen receptor information: 1975-1999. *JNCI. Monographs*, pages 7–15, 2006. (Cited on page 80.)

D. M. Maslove, T. Podchiyska, and H. J. Lowe. Discretization of continuous features in clinical datasets. *Journal of the American Medical Informatics Association*, 20(3):544–553, 2013. (Cited on page 89.)

J. Melnikow, J. J. Fenton, E. P. Whitlock, D. L. Miglioretti, M. S. Weyrich, et al. Supplemental screening for breast cancer in women with dense breasts: A systematic review for the US Preventive Services Task Force. *Annals of Internal Medicine*, 2016. (Cited on pages 5 and 41.)

C. R. Messina, D. S. Lane, K. Glanz, D. S. West, V. Taylor, W. Frishman, and L. Powell. Relationship of social support and social burden to repeated breast cancer screening in the women's health initiative. *Health Psychology*, 23(6):582, 2004. (Cited on pages 32 and 61.)

G. E. Monahan. State of the art — A survey of partially observable Markov decision processes: Theory, models, and algorithms. *Management Science*, 28(1):1–16, 1982. (Cited on pages 11, 12, 51, 64, 76, and 112.)

G. J. Morris, S. Naidu, A. K. Topham, F. Guiles, Y. Xu, et al. Differences in breast carcinoma characteristics in newly diagnosed african–american and caucasian patients. *Cancer*, 110(4): 876–884, 2007. (Cited on page 14.)

NCI. Cancer intervention and surveillance modeling network, 2013. Available from http://cisnet.cancer.gov/. Accessed on March 13, 2013. (Cited on page 12.)

P. J. Neumann and M. C. Weinstein. Legislating against use of cost-effectiveness information. *New England Journal of Medicine*, 363(16):1495–1497, 2010. (Cited on page 51.)

L. A. Newman. Breast cancer in african-american women. *The Oncologist*, 10(1):1–14, 2005. (Cited on pages 8 and 81.)

OECD. *Cancer Care: Assuring Quality to Improve Survival*. OECD Publishing, Paris, France, 2013. (Cited on page 43.)

C. A. Parise and V. Caggiano. Breast cancer survival defined by the ER/PR/HER2 subtypes and a surrogate classification according to tumor grade and immunohistochemical biomarkers. *Journal of Cancer Epidemiology*, 2014, 2014. (Cited on page 80.)

N. H. G. M. Peters, I. H. M. Borel Rinkes, N. P. A. Zuithoff, W. P. T. M. Mali, K. G. M. Moons, and P. H. M. Peeters. Meta-analysis of mr imaging in the diagnosis of breast lesions 1. *Radiology*, 246(1):116–124, 2008. (Cited on page 32.)

R. Pirracchio, M. L. Petersen, M. Carone, M. R. Rigon, S. Chevret, and M. J. van der Laan. Mortality prediction in intensive care units with the super icu learner algorithm (sicula): a population-based study. *The Lancet Respiratory Medicine*, 3(1):42–52, 2015. (Cited on page 122.)

E. D. Pisano, C. Gatsonis, E. Hendrick, M. Yaffe, J. K. Baum, et al. Diagnostic performance of digital versus film mammography for breast-cancer screening. *New England Journal of Medicine*, 353(17):1773–1783, 2005. (Cited on page 4.)

M. L. Puterman. *Markov decision processes: Discrete stochastic dynamic programming*. John Wiley & Sons, 2014. (Cited on page 117.)

N. Reddy, T. Ninan, L. Tabar, and T. Bevers. The results of a breast cancer screening camp at a district level in rural India. *Asian Pacific Journal of Cancer Prevention*, 13:6067–72, 2012. (Cited on page 6.)

S. Rose. Mortality risk score prediction in an elderly population using machine learning. *American Journal of Epidemiology*, 177(5):443–452, 2013. (Cited on page 122.)

S. Russell and P. Norvig. *Artificial Intelligence: A modern approach*. Pearson Education, 2nd edition, 2003. (Cited on page 85.)

B. Sandikci, L. M. Maillart, A. J. Schaefer, and M. S. Roberts. Alleviating the patient's price of privacy through a partially observable waiting list. *Management Science*, 59(8):1836–1854, 2013. (Cited on pages 25, 28, and 52.)

M. E. Schneider. California cuts coverage for mammography. *Family Practice News*, 40(1):5–5, 2010. (Cited on page 6.)

SEER. SEER*Stat Database: Incidence—SEER 9 Regs Resarch Data. Nov 2008 Sub (1973–2006), Linked to County Attributes, Total U.S., 1969–2006 Counties, National Cancer Institute, DCCPS, Surveillance Research Program, Cancer Statistics Branch, released April 2009, based on the November 2008 submission., 2008a. (Cited on page 94.)

SEER. SEER*Stat Database: Mortality— All COD. Aggregated with State, Total U.S. (1969-2006) National Cancer Institute, DCCPS, Surveillance Research Program, Cancer Statistics Branch, released May 2009., 2008b. (Cited on page 94.)

B. Settles. Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 6(1): 1–114, 2012. (Cited on page 87.)

V. L. Shavers and M. L. Brown. Racial and ethnic disparities in the receipt of cancer treatment. *JNCI*, 94(5):334–357, 2002. (Cited on page 15.)

R. Siegel, J. Ma, Z. Zou, and A. Jemal. Cancer statistics, 2014. *CA: A Cancer Journal for Clinicians*, 64(1):9–29, 2014. (Cited on page 61.)

R. L. Siegel, K. D. Miller, and A. Jemal. Cancer statistics, 2015. *CA: A Cancer Journal for Clinicians*, 65(1):5–29, 2015. (Cited on page 1.)

R. D. Smallwood and E. J. Sondik. The optimal control of partially observable Markov processes over a finite horizon. *Operations Research*, 21(5):1071–1088, 1973. (Cited on pages 11, 48, 51, and 111.)

R. D. Smallwood, E. J. Sondik, and F. L. Offensend. Toward an integrated methodology for the analysis of health-care systems. *Operations Research*, 19(6):1300–1322, 1971. (Cited on page 11.)

C. Smigal, A. Jemal, E. Ward, V. Cokkinides, R. Smith, H. L. Howe, and M. J. Thun. Trends in breast cancer by race and ethnicity: update 2006. *CA: A Cancer Journal for Clinicians*, 56(3): 168–183, 2006. (Cited on pages 7 and 77.)

R. Smith-Bindman, D. L. Miglioretti, N. Lurie, L. Abraham, R. B. Barbash, et al. Does utilization of screening mammography explain racial and ethnic differences in breast cancer? *Annals of Internal Medicine*, 144(8):541–553, 2006. (Cited on page 14.)

R. Smith-Bindman, D. L. Miglioretti, R. Rosenberg, R. J. Reid, S. H. Taplin, B. M. Geller, and K. Kerlikowske. Physician workload in mammography. *American Journal of Roentgenology*, 190 (2):526–532, 2008. (Cited on page 6.)

E. J. Sondik and R. Mendelssohn. Information seeking in Markov decision processes. *Southwest Fisheries Center Administrative Report H-79-13, National Marine Fisheries Service, Honolulu, Hawaii*, 1979. (Cited on page 12.)

B. L. Sprague, N. K. Stout, C. Schechter, N. T. Van Ravesteyn, M. Cevik, et al. Benefits, harms, and cost-effectiveness of supplemental ultrasonography screening for women with dense breasts. *Annals of Internal Medicine*, 2014. (Cited on pages 5, 32, 41, and 70.)

N. K. Stout, S. J. Lee, C. B. Schechter, K. Kerlikowske, O. Alagoz, et al. Benefits, harms, and costs for breast cancer screening after us implementation of digital mammography. *JNCI*, 106(6): dju092, 2014. (Cited on pages 4 and 35.)

Susan G. Komen Foundation. Susan G. Komen for the cure and advocacy alliance. State or local actions, 2011. Available from http://www.komenadvocacy.org/statelocalactions.aspx?id=90. Accessed on April 28, 2011. (Cited on page 7.)

J. T. Treharne and C. R. Sox. Adaptive inventory control for nonstationary demand and partial information. *Management Science*, 48(5):607–624, 2002. (Cited on page 11.)

A. B. Tsybakov. Optimal aggregation of classifiers in statistical learning. *Annals of Statistics*, pages 135–166, 2004. (Cited on page 122.)

USPSTF. Screening for breast cancer: Us preventive services task force recommendation statement. *Annals of Internal Medicine*, 151(10):716, 2009. (Cited on page 2.)

N. T. van Ravesteyn, C. B. Schechter, A. M. Near, E. A. M. Heijnsdijk, M. A. Stoto, G. Draisma, H. J. de Koning, and J. S. Mandelblatt. Race-specific impact of natural history, mammography screening, and adjuvant treatment on breast cancer mortality rates in the united states. *Cancer Epidemiology Biomarkers & Prevention*, 20(1):112–122, 2011. (Cited on pages 15, 93, and 94.)

B. Vastag. Breast cancer racial gap examined. *JAMA*, 290(14):1838–1842, 2003. (Cited on page 8.)

V. Velanovich. Immediate biopsy versus observation for abnormal findings on mammograms: an analysis of potential outcomes and costs. *The American Journal of Surgery*, 170(4):327–332, 1995. (Cited on pages 34 and 62.)

F. Vicini, P. Jones, A. Rivers, M. Wallace, C. Mitchell, L. Kestin, I. Jaiyesimi, N. Dekhne, and A. Martinez. Differences in disease presentation, management techniques, treatment outcome, and toxicities in african-american women with early stage breast cancer treated with breast-conserving therapy. *Cancer*, 116(14):3485–3492, 2010. (Cited on page 81.)

A. Vlachos. A stopping criterion for active learning. *Computer Speech & Language*, 22(3):295–312, 2008. (Cited on page 88.)

A. T. Wang, C. M. Vachon, K. R. Brandt, and K. Ghosh. Breast density and breast cancer risk: a practical review. In *Mayo Clinic Proceedings*, volume 89, pages 548–557. Elsevier, 2014. (Cited on page 5.)

M. Wegkamp. Model selection in nonparametric regression. *Annals of Statistics*, pages 252–273, 2003. (Cited on page 122.)

C. R. Weisstock, R. Rajapakshe, C. Bitgood, S. McAvoy, P. B. Gordon, A. J. Coldman, B. A. Parker, and C. Wilson. Assessing the breast cancer risk distribution for women undergoing screening in british columbia. *Cancer Prevention Research*, 6(10):1084–1092, 2013. (Cited on page 62.)

WHO. WHO Breast cancer: Prevention and control, 2015. Available from http://www.who.int/cancer/detection/breastcancer/en/index1.html. Accessed on February 25, 2015. (Cited on page 6.)

B. E. Wojcik, M. K. Spinks, and S. A. Optenberg. Breast carcinoma survival analysis for african american and white women in an equal-access health care system. *Cancer*, 82(7):1310–1318, 1998. (Cited on pages 8 and 14.)

B. C. Yankaskas and K. S. Gill. Diagnostic mammography performance and race. *Cancer*, 104 (12):2671–2681, 2005. (Cited on page 81.)

J. Zhang, B. T. Denton, H. Balasubramanian, N. D. Shah, and B. A. Inman. Optimization of PSA-Based screening decisions for prostate cancer detection. *Manufacturing & Service Operations Management*, 14:529–547, 2012. (Cited on pages 11, 13, and 25.)

A APPENDIX FOR CHAPTER 3

A.1 Relative breast cancer risk by breast density

density	Age < 65	Age ≥ 65
d_0	0.48	0.66
d_1	1.00	1.00
d_2	1.55	1.39
d_3	2.01	1.45

Table A.1: Relative breast cancer risk associated with each density category.

A.2 Observation probabilities

h\o	(η_0, ξ^-)	(η_0, ξ^+)	(η_1, ξ^-)	(η_1, ξ^+)	(η_2, ξ^-)	(η_2, ξ^+)	(η_3, ξ^-)	(η_3, ξ^+)	(η_{no}, ξ^-)	(η_{no}, ξ^+)
$s_0 = (d_0,\ h_0)$	$\nu_a^{d_0}$	$1-\nu_a^{d_0}$	0	0	0	0	0	0	0	0
$s_1 = (d_0,\ h_1)$	$1-\tau_a^{s_1}$	$\tau_a^{s_1}$	0	0	0	0	0	0	0	0
$s_2 = (d_0,\ h_2)$	$1-\tau_a^{s_2}$	$\tau_a^{s_2}$	0	0	0	0	0	0	0	0
$s_3 = (d_1,\ h_0)$	0	0	$\nu_a^{d_1}$	$1-\nu_a^{d_1}$	0	0	0	0	0	0
$s_4 = (d_1,\ h_1)$	0	0	$1-\tau_a^{s_4}$	$\tau_a^{s_4}$	0	0	0	0	0	0
$s_5 = (d_1,\ h_2)$	0	0	$1-\tau_a^{s_5}$	$\tau_a^{s_5}$	0	0	0	0	0	0
$s_6 = (d_2,\ h_0)$	0	0	0	0	$\nu_a^{d_2}$	$1-\nu_a^{d_2}$	0	0	0	0
$s_7 = (d_2,\ h_1)$	0	0	0	0	$1-\tau_a^{s_7}$	$\tau_a^{s_7}$	0	0	0	0
$s_8 = (d_2,\ h_2)$	0	0	0	0	$1-\tau_a^{s_8}$	$\tau_a^{s_8}$	0	0	0	0
$s_9 = (d_3,\ h_0)$	0	0	0	0	0	0	$\nu_a^{d_3}$	$1-\nu_a^{d_3}$	0	0
$s_{10} = (d_3,\ h_1)$	0	0	0	0	0	0	$1-\tau_a^{s_{10}}$	$\tau_a^{s_{10}}$	0	0
$s_{11} = (d_3,\ h_2)$	0	0	0	0	0	0	$1-\tau_a^{s_{11}}$	$\tau_a^{s_{11}}$	0	0

Table A.2: Observation probabilities $f^a(\eta, \xi|(d, h)),\ a \in A^{Scr}$.

h\o	(η_0, ξ^-)	(η_0, ξ^+)	(η_1, ξ^-)	(η_1, ξ^+)	(η_2, ξ^-)	(η_2, ξ^+)	(η_3, ξ^-)	(η_3, ξ^+)	(η_{no}, ξ^-)	(η_{no}, ξ^+)
$s_0 = (d_0,\ h_0)$	0	0	0	0	0	0	0	0	$\nu_a^{d_0}$	$1 - \nu_a^{d_0}$
$s_1 = (d_0,\ h_1)$	0	0	0	0	0	0	0	0	$1 - \tau_a^{s_1}$	$\tau_a^{s_1}$
$s_2 = (d_0,\ h_2)$	0	0	0	0	0	0	0	0	$1 - \tau_a^{s_2}$	$\tau_a^{s_2}$
$s_3 = (d_1,\ h_0)$	0	0	0	0	0	0	0	0	$\nu_a^{d_1}$	$1 - \nu_a^{d_1}$
$s_4 = (d_1,\ h_1)$	0	0	0	0	0	0	0	0	$1 - \tau_a^{s_4}$	$\tau_a^{s_4}$
$s_5 = (d_1,\ h_2)$	0	0	0	0	0	0	0	0	$1 - \tau_a^{s_5}$	$\tau_a^{s_5}$
$s_6 = (d_2,\ h_0)$	0	0	0	0	0	0	0	0	$\nu_a^{d_2}$	$1 - \nu_a^{d_2}$
$s_7 = (d_2,\ h_1)$	0	0	0	0	0	0	0	0	$1 - \tau_a^{s_7}$	$\tau_a^{s_7}$
$s_8 = (d_2,\ h_2)$	0	0	0	0	0	0	0	0	$1 - \tau_a^{s_8}$	$\tau_a^{s_8}$
$s_9 = (d_3,\ h_0)$	0	0	0	0	0	0	0	0	$\nu_a^{d_3}$	$1 - \nu_a^{d_3}$
$s_{10} = (d_3,\ h_1)$	0	0	0	0	0	0	0	0	$1 - \tau_a^{s_{10}}$	$\tau_a^{s_{10}}$
$s_{11} = (d_3,\ h_2)$	0	0	0	0	0	0	0	0	$1 - \tau_a^{s_{11}}$	$\tau_a^{s_{11}}$

Table A.3: Observation probabilities $f^W(\eta, \xi|(d, h))$.

A.3 Performance of the proposed solution approach

Lovejoy [1991b] proposes a lower-bounding scheme using piecewise linear convex representation of the value functions. By using a more general value function definition (i.e., not specific to breast cancer screening problem), we can obtain a piece-wise linear convex representation as follows:

$$v_t^*(\pi) = \max\{\pi\alpha : \alpha \in \Gamma_t\} = \max_{a \in A} \left\{ \sum_{i \in \bar{\mathcal{S}}} \pi(i) \left[w_t(i) + \sum_{j \in \bar{\mathcal{S}}} \sum_{k \in \Omega} p_t^a(j|i) f_t^a(k|i) \alpha^{\iota(\pi,a,k)}(j) \right] \right\} \quad \text{(A.1)}$$

where Γ_t is the set of α-vectors that represent the coefficients of one of the linear pieces of v_t^* and $\iota(\pi, a, k)$ is the index of $\alpha \in \Gamma_{t+1}$ that maximizes $\sum_{i \in \bar{\mathcal{S}}} \sum_{j \in \bar{\mathcal{S}}} \sum_{k \in \Omega} \pi(i) p_t^a(j|i,k) f_t^a(k|i) \alpha(j)$. Note that the inner bracketed term in (A.1) evaluated for the optimal action is a *subgradient* of v_t^* at π, denoted by $\alpha(\pi)$, and the set $\Gamma_t = \{\alpha(\pi) : \pi \in \Pi(\bar{\mathcal{S}})\}$ is sufficient to completely represent v_t^* [Smallwood and Sondik, 1973]. Lovejoy [1991b] notes that, given G be any set of points in $\Pi(\bar{\mathcal{S}})$, and $\hat{\Gamma}_t = \{\alpha(\pi) : \pi \in G\}$, $\hat{v}_t(\pi) = \max\{\pi\alpha : \alpha \in \hat{\Gamma}_t\}$ provides a lower bound on the optimal value function, i.e., $\hat{v}_t(\pi) \leq v_t^*(\pi),\ \forall \pi \in \Pi(\bar{\mathcal{S}})$. Figure A.1 shows a sample piecewise linear value function for $|\bar{\mathcal{S}}| = 2$ and Figure A.2 shows construction of an approximate value function for the

same value function using five grid points.

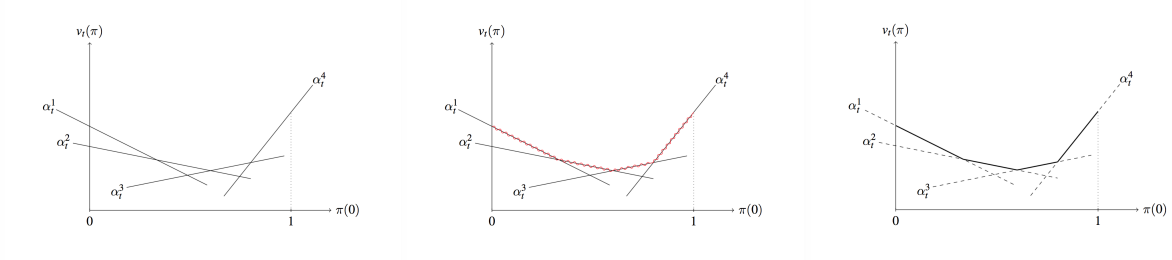

Figure A.1: Piecewise linear value function with four $\alpha-$vectors when $|\bar{\mathcal{S}}| = 2$

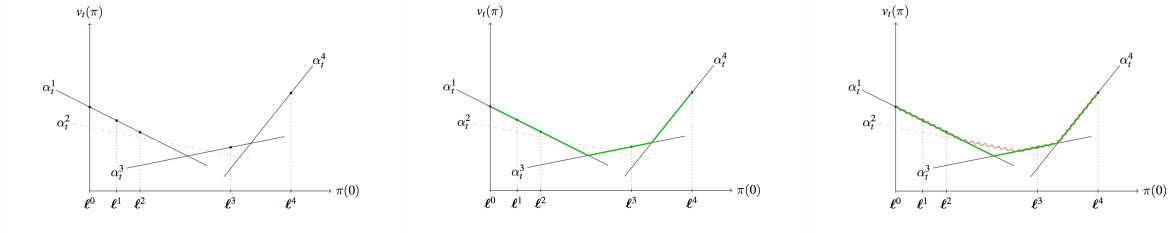

Figure A.2: Approximating the value function with grid points $\boldsymbol{\ell}^0, \ldots, \boldsymbol{\ell}^4$ when $|\bar{\mathcal{S}}| = 2$

As an alternative to Lovejoy [1991b]'s lower bounding approach, we use Eagle's variant of Monahan's algorithm [Monahan, 1982] to devise a lower bounding scheme. Monahan's algorithm exhaustively enumerates all possible α-vectors for a given decision epoch, then eliminates the ones that are dominated by other vectors. Let $\Gamma_{t+1} = \{\alpha_{t+1}^1, \ldots, \alpha_{t+1}^{|\Gamma_{t+1}|}\}$ be the optimal set of α-vectors at decision epoch $t+1$. All α-vectors at decision epoch t can be enumerated as follows:

$$\alpha_t(i) = w_t(i) + \sum_{j \in \bar{\mathcal{S}}} \sum_{k \in \Omega} p_t^a(j|i)\; f_t^a(k|i)\; \alpha_{t+1}^{m(k)}(j), \quad \forall i \in \bar{\mathcal{S}},\; \forall a \in \mathcal{A},\; \forall m \in \mathcal{M}^{t+1}$$

where $\mathcal{M}^{t+1}$ is a collection of tuples of $\{1, \ldots, |\Gamma_{t+1}|\}$ of size $|\Omega|$ (e.g., for $|\Omega| = 4$ and $|\Gamma_{t+1}| = n$, $\mathcal{M}^{t+1} = \{[1,1,1,1],\ [1,1,1,2], \ldots, [n,n,n,n]\}$). Note that $|\Gamma_t| = |\mathcal{A}|\ |\Gamma_{t+1}|^{|\Omega|}$ and it is usually not feasible to enumerate all α-vectors when $|\Omega| > 2$.

Elimination of the dominated vectors is done in two steps. First, α-vectors are compared

component-wise and α-vectors whose components are completely dominated by other α-vectors are eliminated. That is, $\alpha^j \in \Gamma_t$ is dominated if the following condition holds:

$$\exists\, \alpha^k \text{ s.t. } \alpha^k(s) - \alpha^j(s) \geq \epsilon_1, \quad \forall s \in \bar{\mathcal{S}},$$

where ϵ_1 is a precision value. In the second step, a linear programming model is solved for each of the remaining α-vectors to check whether there exists any $\pi \in \Pi(\bar{\mathcal{S}})$ such that the chosen vector dominates every other vector. Monahan's LP is formulated as follows:

$$\begin{aligned}
&\max \sigma \\
&\text{s.t.} \sum_{s\in\bar{\mathcal{S}}} \pi(s)(\alpha_t^j(s) - \alpha_t^k(s)) \geq \sigma, \quad \forall \alpha^k \in \Gamma^t \setminus \{\alpha^j\}, \\
&\quad \sum_{s\in\bar{\mathcal{S}}} \pi(s) = 1, \\
&\quad \pi(s) \geq 0, \quad \forall s \in \bar{\mathcal{S}}.
\end{aligned}$$

If above LP is infeasible or $\sigma \leq \epsilon_2$, then α^j is removed from Γ_t. We use the precision values ϵ_1 and ϵ_2 to obtain a lower bounding scheme. In particular, we set $\epsilon_1 = -10^{-4}$ and $\epsilon_2 = 10^{-4}$, which may lead to some of the optimal vectors to be eliminated from Γ_t. However, this yields Γ_t with a more tractable size, and provides a lower bound on the optimal solution. Our preliminary analysis indicate that the proposed lower bounding scheme performs comparably to Lovejoy [1991b]'s approach.

Table A.4 shows that grid-based model provides good upper bounds on the optimal solution for several problem instances, as the QALY differences and optimality gaps are very small. We also compare the policies obtained by Monahan's algorithm and grid-based model by simulating the lives of 100,000 women between ages 40 and 80. These results indicate that average number of mammographies, MRIs and ultrasounds (E[M], E[R], and E[U], respectively) recommended are generally similar.

model	patient type	Monohan (LB) QALY	E[M]	E[R]	E[U]	grBased (UB) QALY	E[M]	E[R]	E[U]	QALY difference (months)	gap (%)
pomdp_nss	AR	40.316	21.56	0.00	0.00	40.331	17.44	0.00	0.00	0.18	0.04%
	HR	39.898	33.10	0.00	0.00	39.905	31.21	0.00	0.00	0.08	0.02%
	vHR	38.680	38.02	0.00	0.00	38.680	37.88	0.00	0.00	0.00	0.00%
pomdp	AR	40.317	20.85	0.00	0.24	40.331	17.26	0.00	0.41	0.17	0.03%
	HR	39.899	32.53	0.00	0.51	39.906	30.86	0.00	0.60	0.08	0.02%
	vHR	38.687	38.02	0.00	4.42	38.689	37.68	0.00	3.76	0.01	0.00%

Table A.4: Comparison of Monahan's algorithm with grid-based approximation

We also empirically analyze the impact of the resolution settings on solution quality. Table A.5 summarizes the QALY values for a very-high risk patient (vHR) with unknown breast density at age 40 under different resolution settings for health states (he_ress) and density states (den_ress). We observe marginal QALY gains as the resolution improves. For instance, the difference in QALY values is 0.013 months for y^1 and y^4 resolutions. However, y^4 generates significantly more grid points, which leads to increased difficulty in solving the problem.

he_ress	den_ress	numGrids	QALY	QALY diff. (mo.)
y^1 = (250, 100, 10, 5)	1	508	38.689	-
y^2 = (500, 100, 10, 5)	1	1128	38.689	0.007
y^3 = (1000, 200, 20, 10)	1	4328	38.688	0.011
y^4 = (2000, 200, 20, 10)	1	14012	38.688	0.013

Table A.5: Impact of grid resolution on QALY values of vHR patient with unknown breast density

A.4 Value of modeling breast density and supplemental screenings for average-risk and high-risk patients

	avg. Scr	avg. FP	QALY	QALY gain (months)
noScr	0.00	0.00	40.137	-
tri_50_70	6.99	0.47	40.387	3.00
bi_50_70	10.58	0.71	40.415	3.34
bi_40_74	17.19	1.19	40.515	4.54
an_40_74	32.61	2.35	40.486	4.18
pomdp_nss	13.20	0.91	40.573	5.23
pomdp	13.20	0.91	40.573	5.23

(a) Breast density: d_0

	avg. Scr	avg. FP	QALY	QALY gain (months)
noScr	0.00	0.00	39.667	-
tri_50_70	7.47	0.69	40.150	5.79
bi_50_70	10.98	1.06	40.193	6.30
bi_40_74	17.67	1.80	40.332	7.97
an_40_74	32.90	3.48	40.297	7.55
pomdp_nss	15.11	1.49	40.397	8.75
pomdp	15.13	1.49	40.397	8.75

(b) Breast density: d_1

	avg. Scr	avg. FP	QALY	QALY gain (months)
noScr	0.00	0.00	39.250	-
tri_50_70	7.74	0.79	39.921	8.05
bi_50_70	11.30	1.21	39.982	8.78
bi_40_74	17.97	2.08	40.153	10.83
an_40_74	33.00	4.07	40.130	10.56
pomdp_nss	16.05	1.83	40.228	11.74
pomdp	16.41	1.82	40.229	11.75

(c) Breast density: d_2

	avg. Scr	avg. FP	QALY	QALY gain (months)
noScr	0.00	0.00	38.933	-
tri_50_70	8.16	0.75	39.759	9.91
bi_50_70	11.68	1.18	39.834	10.81
bi_40_74	18.28	1.98	40.040	13.28
an_40_74	33.11	3.81	40.039	13.27
pomdp_nss	19.92	2.12	40.126	14.31
pomdp	20.34	2.13	40.126	14.32

(d) Breast density: d_3

Table A.6: Evaluation of screening strategies for average-risk patients

	avg. Scr	avg. FP	QALY	QALY gain (months)
noScr	0.00	0.00	39.381	-
tri_50_70	7.44	0.45	39.898	6.21
bi_50_70	11.04	0.71	39.968	7.04
bi_40_74	17.74	1.17	40.191	9.72
an_40_74	32.86	2.29	40.197	9.79
pomdp_nss	27.13	1.88	40.294	10.96
pomdp	27.49	1.91	40.296	10.98

(a) Breast density: d_0

	avg. Scr	avg. FP	QALY	QALY gain (months)
noScr	0.00	0.00	38.504	-
tri_50_70	8.37	0.65	39.474	11.64
bi_50_70	11.84	1.00	39.573	12.82
bi_40_74	18.63	1.70	39.878	16.48
an_40_74	33.45	3.35	39.891	16.64
pomdp_nss	25.59	2.46	40.004	18.00
pomdp	25.67	2.48	40.004	18.00

(b) Breast density: d_1

	avg. Scr	avg. FP	QALY	QALY gain (months)
noScr	0.00	0.00	37.746	-
tri_50_70	8.87	0.74	39.071	15.90
bi_50_70	12.33	1.13	39.203	17.48
bi_40_74	19.26	1.98	39.572	21.90
an_40_74	33.69	3.88	39.620	22.48
pomdp_nss	31.88	3.62	39.740	23.93
pomdp	32.49	3.64	39.741	23.94

(c) Breast density: d_2

	avg. Scr	avg. FP	QALY	QALY gain (months)
noScr	0.00	0.00	37.180	-
tri_50_70	9.53	0.70	38.788	19.30
bi_50_70	12.86	1.08	38.943	21.16
bi_40_74	19.68	1.84	39.372	26.31
an_40_74	33.94	3.55	39.456	27.32
pomdp_nss	35.78	3.77	39.582	28.83
pomdp	35.72	3.76	39.582	28.83

(d) Breast density: d_3

Table A.7: Evaluation of screening strategies for high-risk patients

B APPENDIX FOR CHAPTER 4

B.1 MDP Model to Approximate POMDP

We describe the components of the MDP model that is used to approximate the POMDP model as follows:

- Decision epochs: $t \in \{0, 1, \ldots, N\}$
- States: $\boldsymbol{g}_i \in \mathfrak{S} \equiv \mathcal{G}$
- Actions: $a \in \mathcal{A} = \{DN, M\}$
- Transition probabilities: $\mathbb{F}_t(\boldsymbol{g}_k|\boldsymbol{g}_i, a)$
- Rewards:

$$\omega_t(\boldsymbol{g}_i, W) = r_t(\boldsymbol{g}_i, W)$$
$$\omega_t(\boldsymbol{g}_i, M) = r_t(\boldsymbol{g}_i, M) + \sum_{s\in\{1,2\}} \boldsymbol{g}_i(s) z_t(M + |s, M) R_t(s)$$

- Optimality equations:

$$\hat{Q}_t^*(\boldsymbol{g}_i) = \max\left\{\hat{Q}_t^{DN}(\boldsymbol{g}_i),\ \hat{Q}_t^{M}(\boldsymbol{g}_i)\right\}, \quad t < N,\ \boldsymbol{g}_i \in \mathfrak{S},$$

$$\hat{Q}_N^*(\boldsymbol{g}_i) = \sum_{s\in\bar{\mathcal{S}}} \boldsymbol{g}_i(s) R_N(s), \quad \boldsymbol{g}_i \in \mathfrak{S},$$

$$\hat{Q}_t^a(\boldsymbol{g}_i) = \omega_t(\boldsymbol{g}_i, a) + \sum_{\boldsymbol{g}_k\in\mathfrak{S}} \mathbb{F}_t(\boldsymbol{g}_k|\boldsymbol{g}_i, a)\ \hat{Q}_{t+1}^*(\boldsymbol{g}_k).$$

Most commonly used method for solving the recursive value functions for finite horizon MDPs is to use a backward induction type algorithm. However, an LP formulation can also be used to solve the MDP models. In order to solve above optimality equations, we use the

following linear programming formulation:

$$\min \sum_{\boldsymbol{g}_i \in \mathfrak{S}} \delta_i \, \mathscr{Q}_0(\boldsymbol{g}_i)$$

$$\text{s.t.} \quad \mathscr{Q}_t(\boldsymbol{g}_i) \geq \omega_t(\boldsymbol{g}_i, a) + \sum_{\boldsymbol{g}_k \in \mathfrak{S}} \mathbb{F}_t(\boldsymbol{g}_k | \boldsymbol{g}_i, a) \, \mathscr{Q}_{t+1}(\boldsymbol{g}_k), \quad \boldsymbol{g}_i \in \mathfrak{S}, \; t < N,$$

$$\mathscr{Q}_N(\boldsymbol{g}_i) = \sum_{s \in \bar{\mathcal{S}}} \boldsymbol{g}_i(s) R_N(s), \quad \boldsymbol{g}_i \in \mathfrak{S},$$

$$\mathscr{Q}_t(\boldsymbol{g}_i) \;\; free, \quad \boldsymbol{g}_i \in \mathfrak{S}, \; t \leq N,$$

where δ_i specifies an initial distribution over the state space.

B.2 A Counterexample for Obtaining Deterministic Policies from Constrained MDPs

It is easy to see that we can always find a deterministic policy for the unconstrained problem using the grid-based approximation, as we can solve the approximate model using a backward induction algorithm [Puterman, 2014]. However, when we add constraints to the model, we may not always get a deterministic policy as a solution. Model having multiple optimum solutions might be a reason for getting randomized policies, but there could be cases where there are no equivalent optimal deterministic policies for the constrained problem. Example 6.9.1 in Puterman [2014]'s book shows that there is no deterministic equivalent optimal policy for the infinite-horizon constrained MDPs. Below we provide a counter example for the finite horizon MDPs using a toy problem given by Puterman [2014].

Consider a two-state MDP example given in Figure B.1, where we aim to find an optimal policy over $N = 3$ horizon. Formal description of the problem is given as follows:

- States: $S = \{s_1, s_2\}$
- Actions: $A_{s_1} = \{a_{1,1}, a_{1,2}\},\ A_{s_2} = \{a_{2,1}\}$
- Immediate rewards: $r_t(s_1, a_{1,1}) = 5,\ r_t(s_1, a_{1,2}) = 10,\ r_t(s_2, a_{2,1}) = -1$
- Terminal rewards: $r_N(s_1) = 0,\ r_N(s_2) = 0$
- Transition probabilities:

  $p_t(s_1|s_1, a_{1,1}) = 0.5,\ p_t(s_2|s_1, a_{1,1}) = 0.5$

  $p_t(s_1|s_1, a_{1,2}) = 0,\ p_t(s_2|s_1, a_{1,2}) = 1$

  $p_t(s_1|s_2, a_{2,1}) = 0,\ p_t(s_2|s_2, a_{2,1}) = 1$

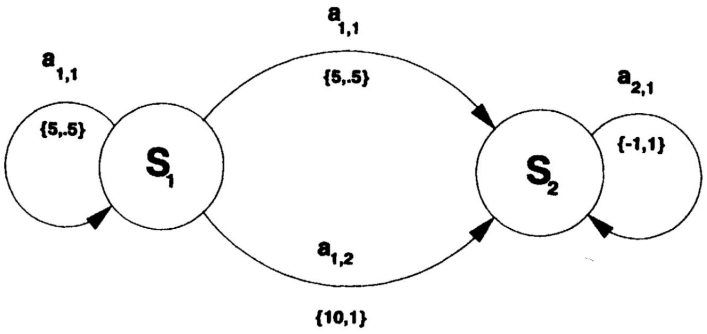

Figure B.1: MDP example

This problem can be solved using a backward induction algorithm. Alternatively, we can use the following linear programming formulation to find an optimal policy for this problem:

$$
\begin{aligned}
\textbf{MDP-dual}\text{: max} \quad & \sum_{t=1}^{N-1} \sum_{s\in S} \sum_{a\in A_s} r_t(s,a)x_t(s,a) + \sum_{s\in S} r_N(s)x_N(s) \\
\text{s.t.} \quad & \sum_{a\in A_s} x_1(s,a) = \alpha(s), \quad \forall s \in S \\
& \sum_{a\in A_s} x_t(s,a) - \sum_{s'\in S} \sum_{a\in A_{s'}} p_t(s|s',a)x_{t-1}(s',a) = 0, \quad \forall t \in \{2,..,N-1\},\ \forall s \in S \\
& x_N(s) - \sum_{s'\in S} \sum_{a\in A_{s'}} p_{N-1}(s|s',a)x_{N-1}(s',a) = 0, \quad \forall s \in S \\
& x_t(s,a) \geq 0,\ \forall s \in S,\ a \in A_s,\ \ x_N(s) \geq 0,\ \forall s \in S
\end{aligned}
$$

where we arbitrarily choose initial probabilities for occupying each state as $\alpha(s_1) = \alpha(s_2) = 0.5$. Solving MDP-dual, we find an optimal objective value of 3.75, and the solution $x_1^*(s_1, a_{1,1}) = 0.5,\ x_1^*(s_1, a_{1,2}) = 0,\ x_1^*(s_2, a_{2,1}) = 0.5, x_2^*(s_1, a_{1,1}) = 0,\ x_2^*(s_1, a_{1,2}) = 0.25,\ x_2^*(s_2, a_{2,1}) = 0.75,\ x_3(s_1) = 0,\ x_3(s_2) = 1$. The optimal policy is $d_1^*(s_1) = a_{1,1},\ d_1^*(s_2) = a_{2,1},\ d_2^*(s_1) = a_{1,2},$ and $d_2^*(s_2) = a_{2,1}$.

Suppose we solve MDP-dual together with the extra constraint

$$\sum_{t=1}^{N-1} \sum_{a\in A_{s_1}} x_t(s_1,a) \leq 0.6.$$

We refer to this new formulation as CMDP-dual. Optimal objective value of CMDP-dual is 3.6 and the solution of the model leads to randomized decision rule $\bar{d}_1^*$ for the first decision epoch, in which

$$q_{\bar{d}_1^*(s_1)}(a_{1,1}) = 0.4,\ \ q_{\bar{d}_1^*(s_1)}(a_{1,2}) = 0.6.$$

That is, if we are in state s_1 at the first decision epoch, we should select action $a_{1,1}$ with probability 0.4 and action $a_{1,2}$ with probability 0.6.

By enforcing deterministic policy constraints as in Section 4.2.1 of the revised paper, we can

obtain the best deterministic policy for the CMDP-dual with an optimal objective value of 3.5. For this example, none of the optimal solutions leads to deterministic policies. Noting that the objective value of best deterministic policy for CMDP-dual is worse than the objective value obtained by the randomized policy (3.6 vs 3.5), these results show that it is not guaranteed to obtain a deterministic policy as a solution of the finite-horizon constrained MDP models. We provide this counter example in the Appendix of the revised manuscript.

B.3 Graph of Stochastic Orderings

Algorithm B.1: Construct graph of stochastic orderings between grid points

Input: $\mathcal{D} = (V(\mathcal{D}), A(\mathcal{D}))$

Sort the grid points in increasing order of first components[†].

$V(\mathcal{D}).add(\boldsymbol{g}_0)$

for $k = 1, \ldots, |\mathcal{G}| - 1$

$V(\mathcal{D}).add(\boldsymbol{g}_k)$

$L = \emptyset$

for $j = 1, \ldots, k-1$

if $\boldsymbol{g}_j \leq_s \boldsymbol{g}_k$

for $i \in L$

if $\boldsymbol{g}_i \leq_s \boldsymbol{g}_j$

$L.remove(i)$

$L.add(j)$

for $\ell \in L$

$A(\mathcal{D}).add(\ell, k)$

[†] Break ties in favor of higher second components.

B.4 Model Validation

We use several previously validated data sources to estimate the parameters of our models. Among the parameters, the most crucial ones are post-cancer life expectancies and health state transition probabilities. In order to estimate the post-cancer life expectancies, we use National Cancer Institute (NCI)'s Surveillance Epidemiology and End Results (SEER) data book, the most comprehensive database on cancer outcomes available in the world. We estimate the transition probabilities, the most critical input of our model, using the University of Wisconsin Breast Cancer Simulation Model (UWBCS).

The UWBCS is a validated microsimulation model of breast cancer epidemiology in United States. The UWBCS was validated using several approaches: First, the UWBCS model projections of incidence, mortality, and stage distribution have been compared to those reported by the SEER program for the period 1975-2010 and it was observed that the UWBCS replicated the patterns of observed US incidence and mortality over time. Second, the UK AGE screening trial was approximated using UWBCS, assuming perfect adherence to invitations for annual screening with 13-year follow-up of women between ages 40 to 49 and it was observed that the model projections closely matched current stage distribution and the AGE trial results. Third, UWBCS was independently cross-validated against incidence and mortality data from Wisconsin cancer registry system, which were not used in the development of the UWBCS. More details about the validation of the UWBCS are available elsewhere [Fryback et al., 2006, Mandelblatt et al., 2016]. We compare the results of our models that are informed by these validated data sources with the findings of Ayer et al. [2012] and observe that we obtain similar results for the unconstrained breast cancer screening problem.

C APPENDIX FOR CHAPTER 5

C.1 Ensembles of ANNs

Ensemble methods construct a set of models from the training data where predictions for the new instances are made by using a weight function that combines individual predictions. Bootstrap aggregating (bagging), boosting and stacking are well known approaches for generating ensembles of the prediction models [Dietterich, 2000]. In addition, some other studies focus on optimal aggregation of the prediction models [Tsybakov, 2004, Wegkamp, 2003]. There are several successful applications of ensemble learning in healthcare literature. In particular, super learning, a generalization of stacking method where another prediction model is employed to combine the individual predictions, has been used for mortality risk prediction in elderly people and mortality prediction in intensive care units [Rose, 2013, Pirracchio et al., 2015].

We observe that using ensembles of ANNs often improve on the performance of a single ANN [Granitto et al., 2005]. Our preliminary experiments with frequently-used ensemble methods such as bagging, stacking, and additive regression indicate that using bagging to obtain an ensemble of ANNs leads to a prediction model with high predictive accuracy. Different than other ensemble methods, bagging forms several replicates of the training data set by sampling with replacement from the original set and a separate learning model is constructed for each of the replicate samples. Bagging may provide substantial gains in accuracy for unstable learning methods such as ANNs where the learning method is significantly affected by the changes in the training set [Breiman, 1996].

C.2 Active Learning Algorithm

Algorithm C.1: Active Learning for Simulation Calibration

Input: Scored instances set L, unscored instances set U, random sample size B'
average number of instances in a cluster k, query batch size B, query strategy $\varphi(\cdot)$

$M = train(L)$ // *learn a model using* L

repeat

$S = \emptyset$

for $b = 1$ **to** kB

$x_b^* = arg\ min_{x \in U \setminus S}\ \varphi(x)$ // *select the lowest scored instance*

$S = S \cup x_b^*$

end

$\{C_1, \ldots, C_B\}$ = **cluster**(S)

for $b = 1$ **to** B

$x_b^* = arg\ min_{x \in C_b}\ \varphi(x)$ // *select the lowest scored instance in cluster* C_b

evaluate(x_b^*) // *evaluate instance to obtain its score*

$L = L \cup x_b^*$

$U = U - x_b^*$

end

$S' =$ **select**(U, B') // *select a random sample of size* B' *from* L

// *evaluate instances to obtain their scores*

for $b' = 1$ **to** B'

evaluate($x_{b'}$)

end

$L = L \cup S'$

$U = U \setminus S'$

$M = train(L)$

until *some stopping criteria;*

C.3 Impact of Training Set Size on Performance of Prediction Models

We demonstrate how bagANN performance varies as a function of the training set size and report the estimates on the additional number of runs required for different threshold values in Table C.1. We compare different training sets consisting of 10,000 to 50,000 instances to demonstrate the effect of the training set size. Again, we use a 60/40 ratio to separate these sets into training and test sets. For this experiment, we sample 30 different training sets for each training set size and provide the average values for the performance measures. As the training set size increases, MAE and RMSE improve overall and for the low-scored instances. However, because we randomly select the initial training sets, there exist discrepancies in the reported results. For instance, while training set size of 50,000 leads to the best MAE and RMSE values for the low-scored instances, MAE and RMSE values for the overall instances are worse compared to the training set size of 40,000. Despite these inconsistencies, our results indicate that, as the training set size increases, bagANN model?s predictive accuracy increases at a higher rate for low-scored instances and the number of additional evaluations decreases considerably. On the other hand, increasing the size of the training set beyond 30,000 instances do not lead to significant performance gains. For example, when the training set size is 30,000, by setting the threshold value $\delta = 15$, we obtain 97% of all acceptable instances by performing 1903 additional simulation evaluations, whereas when the training set size is 40,000 and $\delta = 15$, we obtain 99% of all acceptable instances by evaluating 1936 additional parameter combinations.

C.4 Sensitivity Analysis for Active Learning Parameters

We conduct a one-way sensitivity analysis on initial training set size and batch size parameters of the active learning algorithm. Figure C.1 shows that the active learning algorithm performs better for the initial training sets with 1500-2000 instances. In addition, Figure C.2 shows that

Training Set Size	MAE	RMSE	MAE-low	RMSE-low	ENRuns ($\delta = 10$)	ENRuns ($\delta = 15$)	ENRuns ($\delta = 20$)
10000	2.70	3.71	6.38	7.57	729(39%)	2504(88%)	6624(99%)
20000	2.26	3.05	5.13	6.14	754(64%)	1983(95%)	4443(100%)
30000	2.20	2.94	4.60	5.51	614(76%)	1903(97%)	4086(100%)
40000	2.12	2.85	4.54	5.45	764(81%)	1936(99%)	4136(100%)
50000	2.19	2.88	4.53	5.39	677(84%)	2061(100%)	4079(100%)

Table C.1: Performance of the bagANN for different training set sizes.

although performance of the active learning algorithm is not significantly affected by the choice of the batch size, smaller batch sizes should be more preferable.

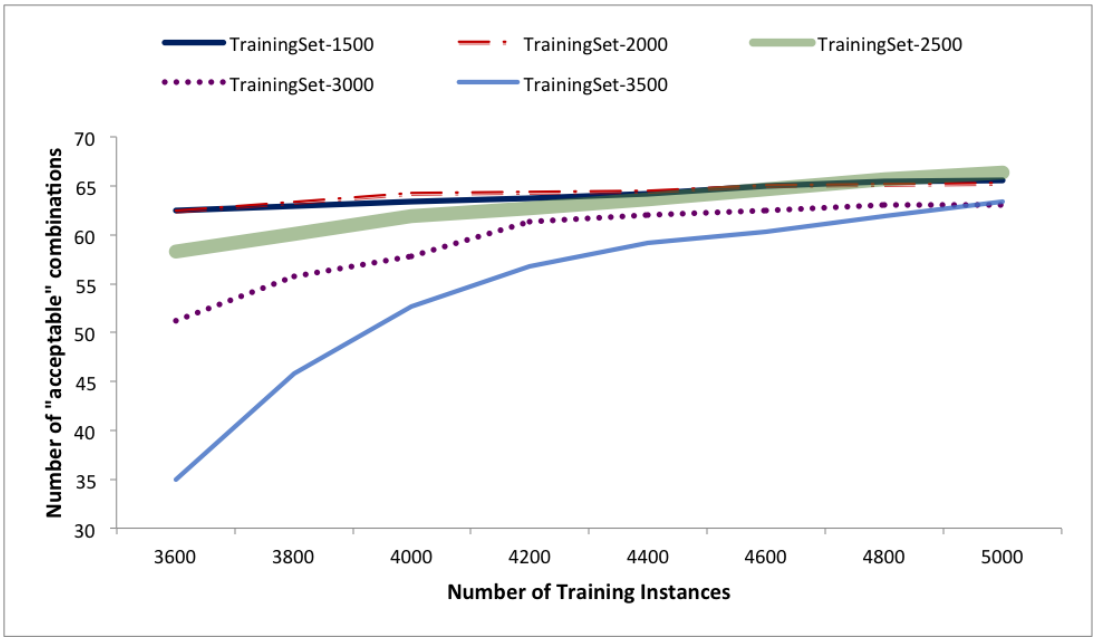

Figure C.1: Learning curves for different initial training set sizes.

In our final experiment, we use the stopping condition based on warm-up period and consecutive number of iterations without new acceptable instances in the training set. We find that, on average, the active learning algorithm stops when the total number of evaluated parameter combinations reaches 5620. These 5620 instances contain all 69 acceptable parameter combina-

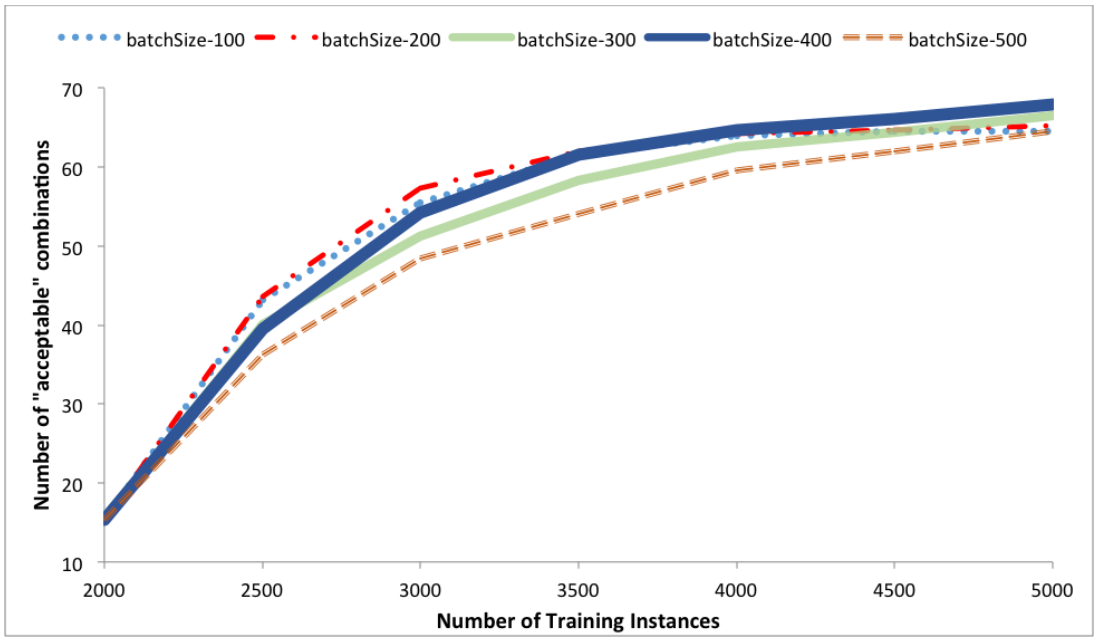

Figure C.2: Learning curves for different batch sizes used in the active learning algorithm.

tions, which implies that evaluating only 1.49% of all 378,000 instances would be sufficient to obtain all acceptable parameter combinations. We also perform a one-way sensitivity analysis on the stopping condition parameters of the active learning algorithm. Table C.2 shows that while increasing the k value leads to more simulation evaluations, it does not lead to a significant improvement in the number of acceptable parameter combinations. Similar analysis on the length of warm-up period shows that our active learning algorithm is not sensitive to the changes in warm-up period. However, using a limit on the length of warm-up period ensures that active learning algorithm does not terminate prematurely.

	$k=2$	$k=3$	$k=4$	$k=5$	$k=6$
avg. # Instances	5340	5620	5820	6200	6640
avg. # Acceptable	65.6	65.6	65.6	66.7	67.9

Table C.2: Performance of the active learning algorithm for different stopping criterion parameters (k).