Optimization Approaches to Identifying Personalized Colorectal Cancer Screening Strategies

By

David Young

A dissertation submitted to the Graduate Faculty of
Auburn University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Auburn, Alabama
August 6th, 2022

Keywords: Cancer Screening; Stochastic Program; Derivative-free Optimization; Colorectal cancer; Preventative screening.

Approved by

Selen Cremaschi, B. Redd & Susan W. Redd Endowed Eminent Scholar Chair Professor
Jin Wang, Walt and Virginia Woltosz Professor
Peter He, Associate Professor
Mark Carpenter, Professor

Abstract

This dissertation investigates optimization approaches for solving a problem coined the colorectal cancer screening problem (CRCSP). This problem aims to identify a screening strategy for testing an asymptomatic population to maximize the benefit of screening through either monetary or societal means. The CRCSP includes two types of uncertainty, exogenous and endogenous uncertainty, where the decision variables impact the probabilities associated with the uncertainty outcomes. The decisions all have discrete values making the problem at best a mixed integer problem. We investigated solving the CRCSP through two different approaches. The first approach, simulation optimization, integrates a microsimulation (MSM) model within a derivative-free optimization (DFO) framework to search for the optimal screening strategy for the simulated population. The second approach uses mathematical programming, specifically stochastic programming (SP), as the framework to model and solve the problem.

To implement the simulation-optimization approach, an MSM of the colorectal cancer progression was first reconstructed, and its outputs were verified using literature data. A comprehensive study of DFO solvers was then performed to identify the solver that is best suited to efficiently and reliably solve combinatorial optimization problems. It was found that the commercial solver TOMLAB/glcSolve, an implementation of the dividing hyper-rectangles (DIRECT) algorithm, performs best for combinatorial problems with a low number of decision variables. The next best solver is the derivative-free line search (DLF) algorithm. This solver shows a similar performance to that of glcSolve. As the number of decision variables increases, however, DFL outperforms glcSolve. Using glcSolve, an optimal screening strategy was identified that provided a 31% improvement in quality-adjusted life-years (QALY) gained while only increasing the overall costs by 5 % compared to the currently recommended strategy. The DFO

framework was then used to analyze the impact of different simulation assumptions and parameter values on the overall optimal screening strategy and the optimal solution. The study showed that the optimum screening strategy identified was most sensitive to the changes in the relative risk (RR) associated with the transition probability of colorectal cancer (CRC), followed by the compliance modeling and willingness to pay ratio (WTP). Changes in the RR changed screening frequency and the screening starting and ending ages. Compliance modeling mainly impacted the screening modality. When WTP changed, the screening start age was affected.

The CRCSP was modeled in two different ways using an SP framework. The first way was to represent the CRCSP as a two-stage SP (TSSP), where the type I endogenous uncertainty caused the problem to be a mixed integer non-linear program (MINLP). Two different direct linearization procedures were then applied to the MINLP, and the resulting models and their solution times were assessed, finding a size versus solution time trade-off for the two linearization procedures. The solution of times of both models were found to be significantly dependent on the size of the scenario set, leading to an investigation of efficient scenario set construction methodologies to best represent the rare-event region of the uncertainty space. The results revealed that the distance-based clustering methods, k-means and x-means, provided very stable and accurate scenario sets compared to a number of sampling schemes.

The second approach to the CRCSP represented the problem as a multi-stage SP (MSSP). The uncertainty was modeled differently than in the TSSP to maintain computational tractability. However, this change led to not having a closed-form expression to calculate the benefits of screening. This problem was overcome by integrating a machine learning (ML) model within the MSSP formulation to estimate the screening benefits. The data to train the ML model was gathered

from the MSM reconstructed from literature. The resulting formulation allowed for a computationally tractable and easily solvable MSSP for the CRCSP.

Acknowledgements

I would first like to thank my advisor, Dr. Selen Cremaschi, for the invaluable guidance, support, and encouragement she provided for me through the entirety of my time at Auburn University. Without her, I would not be where I am today professionally and personally.

I would also like to extend my gratitude to Dr. Jin Wang, Dr. Peter He, and Dr. Mark Carpenter for their feedback in my work to improve my dissertation. I would also like to thank Dr. Aleksandr Vinel who generously donated his time to be the reader for my dissertation. Additionally, I would like to thank my colleagues at Auburn University for providing an incredible research environment.

Lastly, I want to express my deepest gratitude to my parents and sister for providing me with as much love and support as I could ask for, and a large thank you to all of my friends that helped me stay grounded and sane through the entire time of my Ph.D. work.

Table of Contents

List of Figures

List of Tables

CHAPTER 1
INTRODUCTION

Colorectal cancer (CRC) is the third most common and the second most deadly form of cancer worldwide, according to the World Health Organization[1]. Deaths associated with CRC are closely linked to how far cancer has progressed within the individual, with the late stage 5-year survival rate for CRC being around 12 %, while the early-stage 5-year survival rate is upwards of 90 %. This attribute indicates that the early detection of CRC and its precursors can be an effective mitigation strategy to reduce the overall impact of CRC on society. Early detection can be achieved for CRC, and most forms of cancer, by screening an asymptomatic individual using some screening test(s) throughout an individual's life.

The identification of the best, personalized screening strategy can be formulated as an optimization problem. Here, the objective is to reduce the overall health and economic impact of CRC on society. The optimal decisions are the best screening strategies identified based on the individual's risk factors. This optimization problem will be referred to as the CRC screening planning problem (CRCSP) for the rest of the document. A typical quantification metric used in literature to assess health impacts is quality-adjusted life years gained due to screening[2]. Quality-adjusted life years are the years an individual lives discounted by the health state that individual is in for that year. The usual measure of economic burden is the total cost associated with CRC treatment and screening.

A screening strategy is defined as what screening test(s) to perform and at what age(s) to perform these tests. Screening tests for CRC can be divided into three major categories: endoscopy, fecal chemical/DNA tests, and virtual endoscopy[3]. The currently recommended screening strategies have been developed for the general population, only considering the familial history

and race as risk factors[4]. However, specific lifestyle-related risk factors, such as obesity and physical inactivity, have been shown to be strongly linked to the development of CRC, more so than in any other form of cancer[5]. The link between the specific lifestyle-related risk factors and CRC development suggests that personalized screening strategies, which consider these risk factors, may reduce the economic and health impacts of CRC compared to a one-size-fits-all approach to screening.

Most studies that address the CRC screening planning problem utilize a what-if analysis approach by postulating several screening strategies and evaluating the economic and health impacts of these strategies using CRC microsimulations. In these studies, the researchers define a predetermined set of variables, e.g., the screening start ages, the screening test frequencies, the screening test type, and the screening ending ages, and evaluate some or all of the unique combinations of them[2,6,7]. Though easier-to-implement and quicker-to-solve than rigorous optimization problems, these analyses likely lead to sub-optimal solutions due to the incomplete evaluation of all potential screening strategies.

There are many uncertainties associated with CRC progression and the CRCSP. All existing models are only estimations of the progression of CRC, as a complete understanding of the progression of the disease is not known[8–11]. All data used to estimate the parameters of these models are based on a single point in time for each patient, as ethically, doctors must remove the potentially harmful lesions if detected, making dynamic monitoring of the disease impossible. When screening, there is no guarantee that a screening test will detect the pre-cancerous lesion, or the test may detect a lesion when there was none in the first place. These screening outcomes may be observed due to technical error associated with the test sensitivity and specificity or human error in diagnosing these lesions at the time of screening[12]. In addition, individuals do not always

show up at the recommended ages for screening, introducing more uncertainty when trying to develop a screening strategy.

1.1 Objectives

This dissertation will:

(1) Determine optimal screening strategies for CRC using two different approaches: simulation optimization and stochastic programming.

(2) Compare the effects of the optimization approaches on the optimal solutions.

1.2 Organization

Chapter 2 provides a literature review of screening studies for CRC and stochastic programs with decision-dependent uncertainty. It also provides background information on methodologies used in Chapters 3 and 4. The simulation-optimization approach for solving the CRCSP is presented in Chapter 3. Section 3.1 describes the MSM used to simulate the progression of and the process of screening for CRC, as well as the verification study for the MSM. In Section 3.2, a study determining the most efficient derivative-free optimization (DFO) algorithm for combinatorial problems is presented and applied to the CRC MSM. A sensitivity analysis of the MSM is presented in Section 3.3, where different model parameters and assumptions are varied to analyze their impact on the optimal screening strategy. Finally, the results of the simulation optimization are presented in Section 3.4. Chapter 4 presents two stochastic programming (SP) models to model and solve the CRCSP. A two-stage stochastic programming (TSSP) model is presented in Section 4.1, where the model formulation is outlined in Section 4.1.1, and a study to determine efficient scenario construction methodologies is given in Section 4.1.2. As an extension to the TSSP, Section 4.2 presents a data-driven multi-stage stochastic programming model for

solving the CRCSP, and the results of the SP approach are given in Section 4.3. Finally, the overall conclusions and future work are presented in Chapter 5.

CHAPTER 2

BACKGROUND

2.1 Simulations and Colorectal Cancer

Microsimulation models (MSMs) are used as tools to study stochastic processes by a wide variety of disciplines, ranging from transportation planning to health care. A microsimulation is broadly defined as a tool used to study social policy by simulating the effects on individual units, whether as small as individual people or as large as a corporation[13]. As computational power increased, so did the ability to investigate and better understand what drives stochastic processes. The healthcare sector has greatly benefited from such models. The use of MSMs has provided insight into the spread of diseases[14] and individual disease progression, areas that were previously only studied after the event happened. The study of cancer progression and the effects of screening and treatments have been a significant area for utilizing microsimulation models.

The MSM models aim to simulate the progression of CRC, or natural history, within an individual. The accepted pathway for modeling CRC is through the adenoma-carcinoma pathway[15], for which a visual representation is given in Figure 2.1. The pathway begins in a healthy individual with a non-visible, precancerous lesion, known as an adenoma. The adenoma usually grows in size. However, there is evidence[16] that suggests a non-visible, or flat, adenoma can progress directly to a cancerous state, though rare. For those adenomas that do not transition immediately, the likelihood of the adenoma becoming cancerous increases with the size of the adenoma. However, not all adenomas transition to a cancerous state; these are called non-progressive adenomas. Once an adenoma transitions to a cancerous state, it progresses through the stages of cancer.

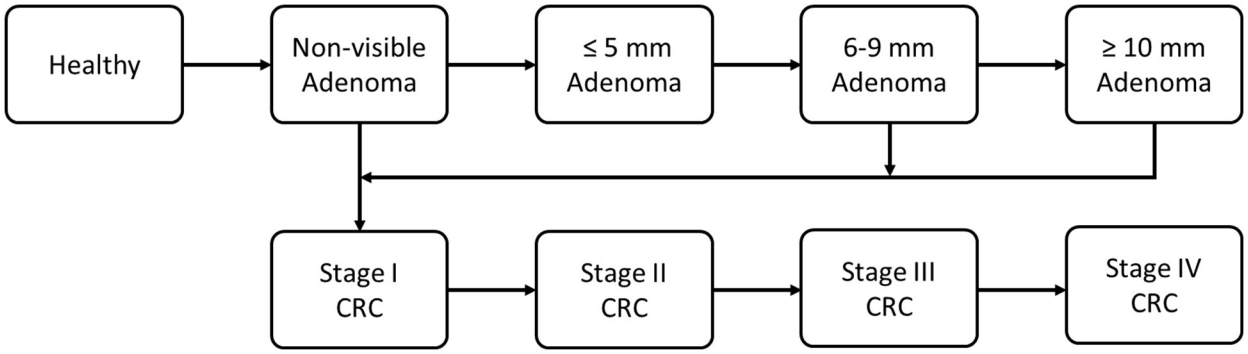

Figure 2.1 Adenoma-carcinoma progression sequence

Due to the uncertain nature of CRC's progression, it is treated as a stochastic process in the natural history section of the MSM. The three most common ways of structuring the natural history to model the stochastic process are using a Markov Chain to step through time[17–24], using a discrete event simulation with the events being various progression stages of the disease[25,26], and finally modeling the whole process as a continuous time model[8–11]. The parameters of the natural history model are used to describe CRC progression, i.e., the probability of developing an adenoma, the risk of an adenoma becoming cancerous, or a case of CRC becoming symptomatic. The model parameters can be estimated with a single value or a distribution. Of the available CRC models, the three most prominent in the US are a part of the US's National Institutes of Health (NIH) Cancer Intervention and Surveillance Modeling Network (CISNET)[27]. These models are Microsimulation Screening Analysis Colorectal Cancer Model (MISCAN-Colon)[12], Simulation Model of Colorectal Cancer(SimCRC)[26], and Colorectal Cancer Simulated Population model for Incidence and Natural history (CRC-SPIN)[8].

The earliest of the three models, MISCAN-Colon was developed in 1999 by Erasmus University Medical Center and Memorial Sloan Kettering Cancer Center. This model follows a Markov Chain framework in which the probabilities to transition from state to state are fit parameters. The risk factors considered in MISCAN-Colon include obesity, diet, physical activity,

multivitamin use, aspirin use, and hormone replacement therapy. These risk factors only affect the development of adenomas and were included based on various clinical studies that analyzed the risk of each factor[28–35]. Additionally, risk factors pertaining to age, sex, and race were included and affect all aspects of CRC progression, not just adenoma development. The calibration method used is based on an adaptation to the Nelder and Mead simplex algorithm to minimize the deviation of simulated statistics, e.g., age-specific adenoma prevalence, CRC incidence by age, or CRC stage, to data taken from clinical trials[36–50] and the US NIH's Surveillance, Epidemiology, and End Results (SEER) database[51]. The minimization problem uses the model parameters as the decision variables.

Developed by the University of Minnesota and Massachusetts General Hospital, SimCRC was the next model to be developed in 2000. The structure of SimCRC follows a discrete event simulation framework, generating the times at which the movement from one health state to the next occurs to simulate the adenoma-carcinoma pathway. The risk factors incorporated into SimCRC are the same as those found within MISCAN-Colon, using the same sources to identify their effect. The method and clinical data used to calibrate model parameters are also similar to MISCAN-Colon. However, the objective of SimCRC's calibration differs from MISCAN-Colon. It is assumed that the calibration datasets (prevalence and number of adenomas by age and sex, location and size/histology of adenomas/carcinomas, and stage and location-specific incidence of CRC by age, sex, and race) follow a multinomial distribution. The log-likelihood values of the parameters for the distributions are then calculated from the data and the modeled output. The objective of the simplex algorithm is to minimize the difference between the two log-likelihood values.

Finally, CRC-SPIN by the RAND Corporation was developed most recently of the three in 2010. The CRC-SPIN is implemented in a continuous time structure, with the progression of CRC as a continuum rather than discrete states as the other two MSM. The model simulates the ages at which adenomas occur within an individual and describes the adenoma growth over time, generating a size at which the adenoma becomes cancerous. Once cancerous, CRC-SPIN then describes the growth of the cancer and ties the stage of CRC with its size. The risk factors incorporated are age and sex, with race incorporated as a new risk factor with revisions to the model. CRC-SPIN authors utilized a Markov Chain Monte Carlo (MCMC) calibration method to estimate the parameter values. With this method, Rutter et al.[52] employ expert opinion to base the prior parameter distributions and incorporate them into the MCMC for calibration. They then use various clinical trials[17,53–57] and the SEER database as their calibration points for output statistics of their Markov chain, i.e., adenoma prevalence, total adenoma per individual, total cases of preclinical cancer based on the size of adenoma, age and sex distribution of CRC cases. Posterior distributions of the parameters are then constructed from the converged MCMC approach, where the acceptance probability of a new point in the approach is a function of the estimated likelihood given the calibration points.

In conjunction with the section modeling the natural progression of CRC, many CRC microsimulations[8,9,24–26,58,10,12,18–23] also include a section within the model for the prevention and early detection of CRC through a screening component. This component allows for a defined screening strategy to be applied to a population, or a specific group of individuals sharing similar characteristics called a cohort, to analyze the effects of the screening. The strategy includes the screening test(s) type(s) performed and the ages at which to screen the population. Usually, the screening section considers screening possible in discrete one-year increments[12,26], where the ages

can be explicitly defined[23] or implicitly defined by specifying the ages to start and end screening and the frequency at which the screening tests should occur[6]. The screening components include different sensitivities and specificities for the various screening tests used based on the size of the adenoma or if the adenoma is cancerous or not, with the sensitivities being based on expert opinion or clinical trials[59–62]. Upon a positive result of the given screening test, appropriate clinical measures are then performed, e.g., a polypectomy is performed on positive screening for a colonoscopy, or a diagnostic colonoscopy is scheduled after a positive stool sample test[3]. The follow-up clinical measures disrupt the natural progression of CRC within the model, creating a modified lifetime for the individuals. The model can then simultaneously compare individual lifetimes without screening and the lifetimes with screening[8] or compare the overall effect of screening on the population/cohort[12,26] under study with a reference population/cohort to see the benefits or drawbacks of the defined screening strategy.

Once a model is calibrated, it is validated to ensure the models accurately represent the progression of CRC and not just replicate the results of the calibrated data. Data from clinical trials or databases that were not used for the model calibration are used to validate the models, e.g., two studies were used for the initial validation of CRC-SPIN [63,64]. The verification experiments are performed by replicating the conditions presented by the data. For CRC-SPIN, MISCAN-Colon, and SimCRC, one validation test[65] used was to simulate a cohort of 65-year-olds and compare the prevalence and size of the adenomas present within the simulated cohorts with the results of a clinical trial[17].

One main focus in screening strategy studies is the cost-effectiveness analysis of different screening strategies[6,7,19,23,26,66–71]. These studies are performed using a what-if-analysis, where the effects the screening strategy has in the prevention or early detection of CRC are compared with a

lifetime with no screening, the natural history, or compared against the effects of another screening strategy. To quantify cost-effectiveness, researchers evaluate the costs attributed to cancer and screening for a population divided by the quality-adjusted life years gained by the population. The quality-adjusted life years gained are calculated by weighting the life years gained by an individual by the stage of cancer at diagnosis, making more advanced stages of CRC a lower quality of life. Cost-effectiveness analysis for screening strategies has been used for assessing both new screening techniques[72] and theoretical screening techniques to understand the most impactful attribute of a screening test or strategy[73]. Zauber *et al.*[72] investigated the cost-effectiveness of using a CT colonography as a screening test compared with the recommended CRC screening strategies. Their evaluation used the three CISNET models to determine the cost-effectiveness of screening strategies using a CT colonography. The results revealed that screening strategies using CT colonography were not cost-effective and strictly worse than the recommended strategies, given the price of the CT colonography. Haug *et al.*[73] investigated the cost-effectiveness of a hypothetical screening test versus FOBT. The hypothetical test had a higher sensitivity for detecting CRC but could not detect adenomas. They used the SimCRC model to determine the cost-effectiveness of the two tests, showing that the test was not cost-effective compared to FOBT. The study concluded that new tests should be able to detect the presence of adenomas to be considered an effective screening test.

The previous studies have been useful in assisting policymakers in determining better screening strategies to suggest to the general population. However, the method used for choosing the screening strategies does not allow for a comprehensive evaluation of potential screening strategies. The work by Knudsen *et al.*[2], which studied one of the largest numbers of potential screening strategies, only considered 204 unique potential strategies. If every unique strategy were

considered based on the range of variables defined within the study, around 1,800 CRC screening strategies would have been evaluated. By only evaluating 11% of the strategies, there is likely a better screening strategy than the ones considered. With a rigorous optimization methodology, all potential screening strategies could be considered rather than only a subset.

2.2 Derivative-free Optimization and Simulation Optimization Applications to Cancer Screening

Derivative-free optimization (DFO), also called black-box optimization or simulation-optimization, is a form of optimization in which the DFO algorithm has no information on the derivative of the problem at hand or the problem itself is too complex to use mathematical optimization[74]. Instead, these algorithms rely on search heuristics based on improving the current best set of decision variables. The general framework for DFO is demonstrated in

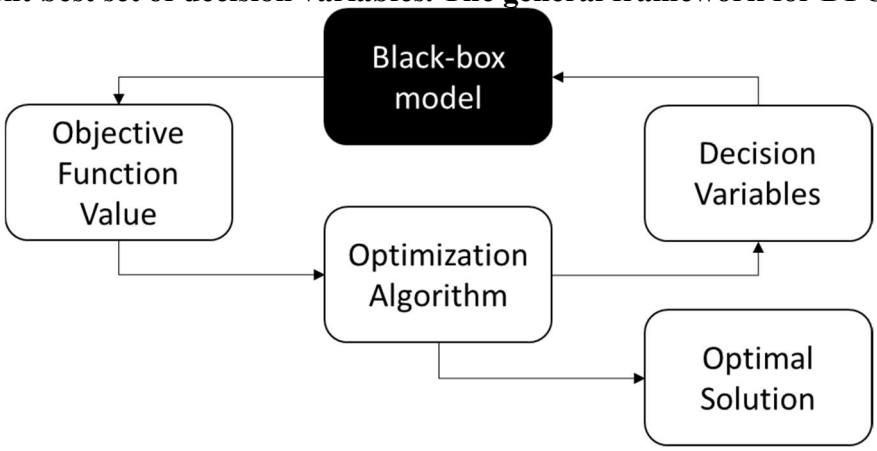

Figure 2.2. Each DFO algorithm begins with a specified initial set of decision variables or uses some sampling method to generate a set. The variable sets are passed to the problem, model, or simulation treated as a black-box model. The only information known by the algorithm is the corresponding objective value for any set of decision variables. With the decision variables evaluated, the objective function values are then passed back to the DFO algorithm, where the algorithm-specific search heuristic begins. The search heuristic uses the input-output pairs of the evaluated variable sets and corresponding objective function values to determine the next set(s) of decision variables to evaluate. This process is iterated until a termination criterion has been

reached. The termination criteria can range from a maximum number of black-box evaluations to a minimum search step size or exceeding a specified wall time.

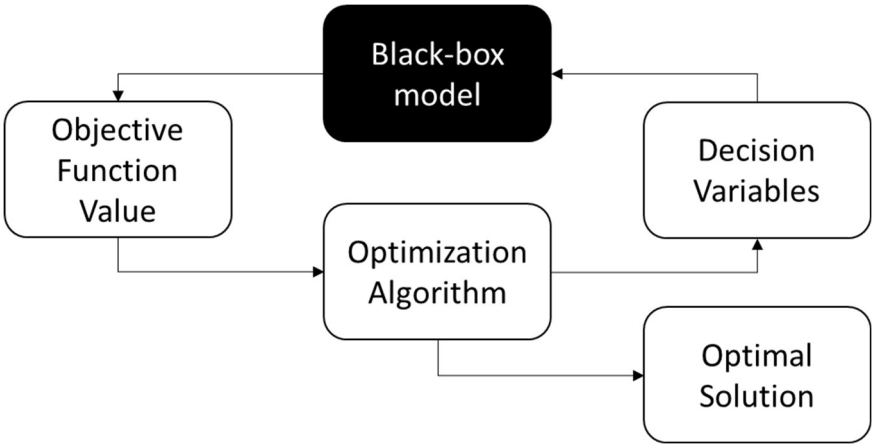

Figure 2.2 Derivative free optimization flowchart

The use of DFO for solving optimization problems tends to be reserved for problems in which the exact formulation of the problem is unknown, part(s) of the problem is represented by a simulation model, or the problem is computationally expensive to evaluate. Therefore, the solutions obtained with DFO are not guaranteed to be the global optima of the problem[75]. Without the guarantee of converging to a globally optimal point, DFO may be considered a worse option than mathematical programming. However, due to its flexibility, DFO can be applied to a wide range of problems, providing at least a locally optimal solution to a problem that may be computationally intractable for a mathematical programming approach.

Simulation optimization has been widely used to determine optimal screening strategies for many forms of cancer, including cervical[76], breast[70], and prostate[77] cancers, to reduce the impact of cancer and due to the availability of natural progression simulation models of these forms of cancer. Bertsimas et al.[77] developed and presented an approach to determine screening strategies that performed well under multiple simulation models. The approach utilized an undefined, iterated local search heuristic that was applied to three simulation models for prostate

cancer from the CISNET consortium. The screening strategy was defined as the starting age, ending age, and frequency of screening and the cut-off values for medical intervention of preostate-specific antigen (PSA) tests as a function of age. Using the approach, they generated a Pareto optimal set of screening strategies, with the two objectives being the maximization of the average quality-adjusted life expectancy (QALE) across the three models and the maximization of the minimum of the three QALE values of the models given the same screening strategy. It was found that all of the strategies generated by their approach performed better than the previous, expert-generated screening strategies. Rauner et al.[70] developed an optimization framework to generate screening strategies for chronic disease that maximized quality adjusted life years gained and minimized cost. Their framework was used to generate the Pareto optimal set of breast cancer screening strategies. This framework looks from the policy maker's point of view to determine the percentage of a population to screen for a disease and how often screenings should occur, including a penalty term if a screening strategy exceeds an annual budget. To solve for the Pareto optimal set, they utilized a multi-objective implementation of Ant Colony Optimization. The study found over 5,000 screening polices that outperformed the current recommended screening policy for breast cancer, suggesting that policy makers should urge a higher compliance rate for older age women. McLay et al.[76] performed a simulation-optimization study to generate dynamic age based screening strategies for cervical cancer. The simulation-optimization was performed using OptQuest for a total number of life time screens ranging from 1 to 22 to determine the optimal ages screening should be performed. The resulting age-based screening strategies showed similar benefits to current, fixed interval, screening strategies while reducing the total number of screens in an individual's lifetime. These three studies all demonstrate that the use of DFO methods are

able to generate screening strategies for cancer that may be better than the current recommendations for the given cancer type.

However, even with the availability of high-fidelity simulation models of CRC, there have not been simulation-optimization studies to investigate optimal CRC screening strategies. In addition, the simulation optimization studies on other cancers selected a DFO algorithm without a study to verify that the selected DFO algorithm reliably solves the underlying problem to optimality. The ten solvers reviewed in Section 2.2.1 were all chosen due to their ability to handle a fully discrete set of decision variables, as the decisions when generating CRC screening strategies have all discrete choices. From the 10 solvers identified, two utilize local search (DFL and NOMAD), six are global search algorithms (DEAP-GA, CMA-ES, SimAnneal, GpyOpt, and TOMLAB/glcSolve and glcDirect), and two employ hybrid search (MIDACO and RBFOpt). Most models use the direct search method, with only three solvers utilizing model-based searches (RBFOpt, GpyOpt, and NOMAD). Additionally, only three solvers follow the same search path at every iteration (TOMLAB/glcSolve, TOMLAB/glcDirect, and DFL), while the remaining seven solvers rely on random number generation for their search path.

2.2.1 Derivative-free Optimization Solver Information

There have been four major reviews[75,78–80] of the available DFO algorithms with the past 10 years, comparing their search strategies and performances over a wide array of test problems. The DFO algorithms can be classified based on the general search strategy used to determine the next set of decision variable values for the black box model to evaluate. These classifications, defined by Boukouvala et al (2016) based on the search domain considered by the algorithm, are local, global, and hybrid search. In addition to search domain, DFO solvers can be classified by the methods in which the next set of decision variables are determined, these can be either direct

search or model based. Model based search uses a surrogate model of the objective function to determine the next trial point, whereas direct search uses the previous set of decision variables to generate a new set of decision variables. They can also be described by if an algorithm's search is deterministic, or if the search requires the use of random numbers. The details of the DFO solvers used in this work are presented in the following sections.

2.2.1.1 Local Search Algorithms

Local search algorithms only generate samples within a specific region around the current best solution. This attribute allows these algorithms to converge more quickly to a solution, but can also make it much harder for them to escape local optima. For non-convex objective functions, this can cause the algorithm to converge to a local optimum instead of finding the global optimum. For this reason, local search algorithms are often used with a multi-start approach to reduce the risk of being trapped in a local minimum[80]. Local search is especially useful in cases where the objective function is convex due to its' ability to quickly and efficiently arrive at the optimum.

2.2.1.1.1 DFL (Derivative-Free Line search) Solver

The core of the DFL solver is a line search algorithm[81]. The algorithm can be used to solve constrained MINLPs and is initialized with a given starting point and a starting step size, and then generates lines in the positive and negative direction from the current best solution. These lines correspond to each of the given decision variables. The algorithm then evaluates points by discretizing the line given the initial step size to check for an improvement in objective function value given the current decision variable. If no improvement is found, the step size on the line is decreased, and the decision variables are revaluated. If there is no improvement for the current best objective function value by the time the algorithm reaches the minimum defined step size, the process is repeated for the remaining decision variables. If an improved solution is found, the step

size is increased, and the algorithm repeats the search from the new point. This process is repeated until a defined stopping criterion is met, i.e., the maximum number of function evaluations, minimum step size for all search directions, or the total number of iterations. For DFL, defining the minimum step size for each discrete variable is necessary as the solver does not automatically set those values for the user.

2.2.1.1.2 NOMAD (Nonlinear Optimization by Mesh Adapted Direct Search) Solver

The underlying algorithm within the NOMAD solver is the mesh adapted direct search (MADS) algorithm[82] to solve constrained MINLPs. After the starting point is chosen, the sample space around the starting point is discretized by a mesh. This mesh creates potential evaluation sites at the intersection points of the mesh and begins with a very coarse grid. With the grid constructed, the algorithm then generates two sets of potential new points, one set for a "poll" step and one set for a "search" step. The algorithm iterates over these two steps adjusting the mesh size and updating the best current solution until the stopping criteria, e.g., the maximum number of black-box model evaluations, minimum mesh size, or convergence on a user defined solution, are met.

The two different steps contribute to the effectiveness of this algorithm. The search step generates a set of new points for the algorithm to consider for the optimal solution. The set of new points is randomly generated from the entire search space using the current grid to define potential points, making this step the global search aspect of the algorithm. These points are then evaluated, and their objective function values are compared to the current best solution. If this search fails to produce a better solution, the mesh size is reduced, and the poll step is called. The poll step generates a set of trial points that are within a set distance of the current best solution. Once again, these points are evaluated, and their objective function values are compared to the current best

solution. When the search or poll step generates a new best solution, the current best point is updated, and the mesh size is increased, creating a coarser grid. If a better solution is not found by either the poll or search step, the mesh becomes finer, and the best solution remains unchanged. After these updates are completed, the algorithm begins a new iteration.

2.2.1.2 Global Search Algorithms

Global search algorithms sample from the entire domain and therefore have a higher potential of escaping a local optimum than local search algorithms. While global optimality is not guaranteed, global search algorithms can usually find the global optimum more often than local search methods for non-convex functions[75]. This larger sample area, however, tends to increase the number of iterations and/or black-box model evaluations it takes for a global search algorithm to converge.

2.2.1.2.1 DEAP – GA (Distributed Evolutionary Algorithms in Python Genetic Algorithm) Package

The DEAP package[83] contains a variety of evolutionary optimization algorithms, such as particle swarm, differential evolution, and genetic algorithm. All of the algorithms in the package are stochastic and follow some of the same general trends. This section focuses on the genetic algorithm (GA) implementation of the DEAP package, as GA, by default, can handle discrete decision variables.

Genetic algorithms[84] are inspired by the evolution of a population and its adaptation via survival of the fittest to solve bounded MINLPs. For GA, the population is a set of different decision variable sets, with the fitness of each set of decision variables within the population being defined based on the objective function. A genetic algorithm is initialized with a randomly generated population of sets of decision variables, or gene sets, with each decision variable being

a gene that is described by a string of bits. Each gene set is then evaluated using the objective function, storing the objective function value as the gene set's fitness value. The algorithm attempts to cull the weaker solutions by populating the next generation with the previous generation's gene sets dependent upon their relative fitness. After the new generation has been populated, random changes across the gene sets are produced through performing genetic methods. The two methods used are called mutation and crossover. The mutation operator takes an individual and randomly swaps the bits within the gene set producing a new set. The crossover operator resembles mating, where two "parent" gene sets produce two "offspring" gene sets. The offspring gene sets are chosen by randomly selecting a position in the parents' genes and swapping the remainder of the set with the other parent's set. The process then iterates until a maximum generation or number of function evaluations is reached, or the population is said to have a converged fitness value defined by some predetermined tolerance. In addition to the ending criteria, the algorithm requires defining the probabilities associated with the occurence of crossover and mutationfor each generation.

2.2.1.2.2 CMA-ES (Covariance Matrix Adaption Evolutionary Strategy) Solver

Stochastic in nature, CMA-ES falls within the category of an evolutionary algorithm, in which the algorithm randomly generates a new set of trial points, or generation, based on a current set of trial points. Initially, this algorithm was implemented to solve bounded nonlinear programs (NLPs)[85]. Similar to the genetic algorithm, the new set of trial points is created through mutation of the current set. This mutation, however, is performed by sampling from a multivariate normal distribution and adding the sample to each current solution in the current generation. The normal distribution has a mean of zero and a covariance that is updated at each generation to improve subsequent generations to find better solutions. Though it was initially created for continuous

problems, one implementation of CMA-ES[86] was constructed in MATLAB to discretize the domain and extend the algorithm to mixed-integer and pure integer problems. For the mixed-integer implementation of CMA-ES, it is necessary to define the "stairwidth" parameter for the discrete or integer decision variables.

2.2.1.2.3 SimAnneal (Simulated Annealing) Solver

The SimAnneal solver searches for the optimal solution using a simulated annealing algorithm[87] for solving bounded MINLPs. This process seeks to mimic the metallurgical process of annealing, where metal is heated and then allowed to cool slowly to relieve internal stresses and strengthen it. To mimic this process, the "temperature" of the algorithm is cooled over time. In this case, the value representing the "temperature" is related to the probability of accepting a solution worse than the current one and is synonymous with the acceptance rate of the worse solutions by the algorithm.

At each iteration, the algorithm takes the current value, or initial value for the first iteration, and generates a new sample based on a user-defined approach to determine a step size to move away from the current solution, or state. The user approach allows the sampling strategy to be as simple as generating uniform random numbers or as complex as incorporating problem-specific knowledge for the sampling procedure. The new sample is then compared with the previous solution. If the new state has a better objective function value than the previous state, it replaces that state. However, an inferior solution is not immediately rejected; instead, it could be accepted as the new state of the system with a probability based on the current temperature of the algorithm. The higher the temperature, the greater the odds that an inferior solution is accepted. Accepting inferior solutions allows the algorithm to step into regions away from the current local minimum. Once the generated solution has either been accepted or rejected, the temperature is decreased, and

the entire process is repeated until the temperature reaches a predetermined value. For SimAnneal, it is necessary to define the total number of steps, and the start and end temperatures.

2.2.1.2.4 GpyOpt (Gaussian Process in Python Optimization) Solver

GpyOpt[88] is an optimization solver based around a Gaussian Process (GP) framework in Python called Gpy. It uses a Bayesian optimization[89] algorithm for solving constrained MINLPs. In Bayesian optimization, the objective function is represented by a surrogate model built using a Gaussian process surrogate model. The surrogate model is updated after every sampled point using Bayesian statistics. An acquisition function is used to determine the next sample point. There are several different types of acquisition functions that GpyOpt can employ, with the default the expected improvement function. The next sample point is determined by maximizing the acquisition function of the GP model. The acquisition function has a balance of exploitation versus exploration, in which the optimal point for the acquisition function for an iteration may be a point to refine the GP model within a poorly fit region or a point that is expected to improve the current best solution. For GpyOpt, the termination criteria are a user-defined maximum number of function evaluations or the identification of the same point by the acquisition function in subsequent iterations.

2.2.1.2.5 TOMLAB/glcSolve and TOMLAB/glcDirect solvers

TOMLAB is a commercial optimization toolbox implemented within MATLAB, with a wide array of different solvers. The solvers used for this dissertation are glcSolve and glcDirect[90]. Both solvers implement a form of the DIRECT (Dividing hyperRECTangles) algorithm[91]. This algorithm partitions the solution space into smaller and smaller hyperrectangles, evaluating these rectangles by the solution found in its midpoint. The algorithm selects the hyperrectangles to investigate further by considering their size and the midpoint function values. Both solvers are

designed to solve constrained MINLPs, with the implementations of DIRECT allowing for both integer and continuous variables. By nature, both solvers do not use any random values.

2.2.1.3 Hybrid Search Algorithms

Hybrid search algorithms rely on neither purely local nor purely global search heuristics. These algorithms use global methods to identify regions where the optimum is most likely to reside. Once these regions have been identified, hybrid algorithms improve their efficiency and decrease their runtime by utilizing local search methods to refine their solutions. This strategy allows hybrid search algorithms to escape local optima while benefiting from faster convergence times.

2.2.1.3.1 MIDACO (Mixed Integer Distributed Ant Colony Optimization) Solver

At its core, MIDACO utilizes a mixed integer Ant Colony Optimization (ACO) algorithm to solve constrained MINLPs. It generates new trial solutions via a multi-kernel Gaussian distribution and handles constraints using a penalty function[92]. In ACO, the algorithm attempts to mimic the actions of ants foraging for food. The algorithm starts with an initial point from which a number of solutions, or ants, are generated. From the generated ants, a set of the best solutions are stored in solution vectors along with their respective penalty functions. The solution vectors are utilized to construct the multi-kernel Gaussian distribution that is sampled to produce the next generation of ants to evaluate. If the best solution from the current generation is better than the worst value saved in the solution vector, it replaces that value, and the distribution is updated. For the penalty function, the oracle penalty method[93] is utilized. This method allows MIDACO to test solutions that slightly violate constraints but seem to be in a promising region by creating slack variables that penalize the objective function.

The algorithm evaluates the change in the average objective function value of the ants at each iteration. Once the change drops below a threshold, the algorithm is considered to be in the final stage. At that point, a set of local solvers are used to fine-tune the answer. These include a pseudo-gradient-based backtracking line search and a mixed integer sequential quadratic programming algorithm[94]. The local search methods of the final step allow MIDACO to converge to a solution much faster than with a pure ACO algorithm. The solver runs until a termination criterion is reached; the criterion can be the maximum number of function evaluations, maximum computation time, or maximum consecutive function evaluations that yield no improvement in the best solution.

2.2.1.3.2 RBFOpt (Radial Basis Function Optimization) Solver

RBFOpt is a framework that utilizes radial basis functions[95] (RBF) to solve bounded MINLPs by generating an RBF surrogate model of the objective function. It then uses this surrogate model with a mixed integer linear program (MILP)/MINLP solver. The default algorithm, Metric Stochastic Response Surface, begins by taking initial points based on the selected initialization sampling strategy, e.g., at the extreme bounds of the search space or Latin Hypercube Sampling[96]. Once the initial points have been chosen, the iteration step starts. The iteration step begins by training, or generating if this is the first iteration step, an RBF model based on the current data set. The user can define what RBF model to use; however, the solver chooses the model with the best fit. Each iteration step is made up of global search and local search, or refinement, phases. The ratio of these phases is controlled by a parameter established by the user, with the default being five global, iterative step, searches to one local, refinement step. The default global search uses a genetic algorithm to generate the trial points using a global RBF representation

of the objective function. A new refinement step begins when a sufficient number of global searches have been completed since the last refinement step.

For the refinement step, the algorithm is seeking to improve upon the current best local solution. This improvement is achieved by using a variation of the trust region method, creating an RBF model to represent the region around the current best solution. An optimization problem is constructed using the RBF model of the region nearest the current best solution. The solution to this optimization problem provides the next sample point. The new point and its objective function value evaluated using the black-box model are added to the solution set. If there is an improvement, this new sample point replaces the point in the set with the worst objective value for both the local and the global RBF models. The radius of the trust region is adjusted based on the previous sample, and the process repeats. Once the refinement step ends, all new points that have been evaluated are added to the main solution set. The iteration and refinement steps continue to alternate until either the maximum number of iterations or wall time has been reached.

2.3 **Stochastic Programming**

Stochastic programming (SP) is a mathematical optimization framework used to determine optimal decisions under uncertainty. The simplest case of SP is known as a two-stage stochastic program (TSSP). In this problem, there are two sets of decisions, the first stage decisions or the here-and-now decisions, and the recourse actions, the corrective actions taken after the uncertainty is realized. The uncertainty is incorporated into the SP framework by constructing scenarios that represent the uncertainty information. The scenarios represent different sets of realizations of the uncertainty. The solution for a TSSP is the first stage decisions that maximize, or minimize, the expected value of the objective function over all scenarios[97].

Two-stage SPs can be extended into multi-stage stochastic programs (MSSPs). This extension allows sequential decisions to be made at multiple stages or time points as uncertainty is gradually realized. After a decision is made for a stage, the uncertainty associated with that stage is realized, creating new scenarios. Figure 2.3 demonstrates a time horizon for a three-stage MSSP. The different time points represent the times at which the uncertainty is realized. At the beginning of the horizon, a decision, the initial decision, is made. At proceeding time points, both recourse actions and decisions can be taken dependent on the current realization of uncertainty[97].

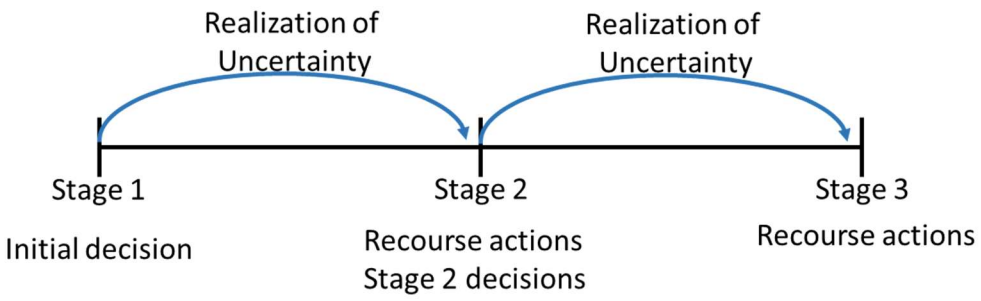

Figure 2.3 Example planning horizon for a three-stage stochastic program

2.3.1 Stochastic Programming and Scenario Construction

A general formulation for a two-stage stochastic program (TSSP) is presented in Eqns. 2.1-2.4[98]. In this problem, the objective is to minimize the value of Z, where $\boldsymbol{x}$ are the first stage decision variables, $\boldsymbol{y}$ are the second stage decision variables, or recourse actions, and ω is the problem uncertainty where it can take on any value within the uncertainty space defined by the probability space (ω, Ω, p), where $p(\omega)$ is the probability measure of outcome ω. The function $f(\boldsymbol{x})$ defines the contributions of the first stage decisions to the objective, $\mathbb{E}_{\omega}\big(Q(\boldsymbol{x}, \omega)\big)$ is the expectation of the function $Q(\boldsymbol{x}, \omega)$ that defines the contributions of the second stage problem to the objective over the uncertainty of the problem, and $h_1(x)$ defines the constraints for the first stage of the problem.

$$\text{Min}\, Z = f(\boldsymbol{x}) + \mathbb{E}_{\omega}\big(Q(\boldsymbol{x},\omega)\big) = f(\boldsymbol{x}) + \int_{\omega\in\Omega} p(\omega)Q(\boldsymbol{x},\omega)d\omega \quad 2.1$$

$$s.t. \quad h_1(\boldsymbol{x}) \leq 0 \quad 2.2$$

$Q(\boldsymbol{x},\omega)$ is determined to be the optimal values of the second-stage optimization problem given in Eqns. 2.3 and 2.4, where $g(\boldsymbol{x},\boldsymbol{y},\omega)$ defines the contribution of the second stage problem to the objective value given the first stage decisions, uncertainty realization, and the recourse actions taken and $h_2(\boldsymbol{x},\boldsymbol{y},\omega)$ defines the constraints of the second stage problem.

$$\min Q(\boldsymbol{x},\omega) = g(\boldsymbol{x},\boldsymbol{y},\omega) \quad 2.3$$

$$s.t. \quad h_2(\boldsymbol{x},\boldsymbol{y}) \leq 0 \quad 2.4$$

The SP problem quickly becomes computationally intractable in this form (Eqns. 2.1 – 2.4) due to the integral over an optimization problem. To remedy this, the problem can be written to approximate the expectation operator, as shown in Eqn. 2.5. It should be noted that in the case of a finite number of realizations of Ω, this is an exact representation of the expectation over the second stage problem. The approximation relies on the introduction of set S, where elements of set $S, s \in S$, are used to index over a finite vector of realizations of Ω, referred to as scenarios, denoted by ω_s.

$$\mathbb{E}_{\omega}\big(Q(\boldsymbol{x},\omega)\big) \approx \sum_{s\in S} p(\omega_s)Q(\boldsymbol{x},\omega_s) \quad 2.5$$

This approximation then allows the TSSP (Eqns. 2.12.4) to be written in, what is called, the deterministic equivalent formulation, Eqns. 2.6-2.8, which is more computationally tractable formulation.

$$\text{Min}\, Z \approx f(\boldsymbol{x}) + \sum_{s\in S} p(\omega_s)g(\boldsymbol{x},\boldsymbol{y}_s,\omega_s) \quad 2.6$$

$$s.t. \quad h_1(\boldsymbol{x}) \leq 0 \quad 2.7$$

$$h_2(\boldsymbol{x}, \boldsymbol{y}_s, \omega_s) \leq 0 \qquad \forall s \in S \qquad 2.8$$

However, when the total number of possible realizations is too high or the domain of Ω is continuous, a set of scenarios must be constructed to represent the uncertainty space properly or a solution approach must be taken to provide a quality approximation of the true solution. One such approach is Sample Average Approximation (SAA)[99]. In SAA, it is assumed that the full, or a sufficiently large, scenario set can be employed for the second stage problem. Then, a subset of the scenarios is generated and used to determine the first stage decisions; those decisions are then fixed and used to solve the second stage problem to estimate the contribution of the second stage decisions to the objective function value. This process is iterated a number of times where the best solution over all iterations is chosen to be the approximated solution.

Several methods have been proposed to generate scenarios efficiently. Høyland and Wallace[100] present a scenario tree construction methodology to generate scenarios for SPs with any number of stages. They base their methodology around matching specified statistical properties of the problem uncertainty. To generate the tree, they formulate an optimization problem to determine the uncertainty realizations, x, and the corresponding probability of occurrence, p. The objective function minimizes the sum of a distance measure, such as the square norm, between the properties of the original distribution/data and the resulting scenario tree, including a weight factor to allow for emphasis of a given statistical property if there are trade-offs between the properties. To expanding the approach to a multi-stage problem, the authors introduce state dependent variables, such as stage specific mean or standard deviation, that track the realizations of uncertainty over time.

Calfa et al.[101] propose a data-driven optimization approach for scenario generation. The work presents an approach to solving the Distribution Matching Problem (DMP). This problem

aims to determine the optimal uncertain parameter realizations and values to best match marginal empirical distribution and the marginal moments of the uncertain parameters. The study discusses two methodologies, one for TSSPs and another for MSSPs, as well as five formulations of the DMP, three NLPs, and two linear programs (LPs). The NLP decision variables include the probabilities of each scenario and the values of the uncertain parameters, with the models differing based on how the moment matching is evaluated, i.e., using the L^1, L^2, or L^∞-norms. The LP decision variables, however, only include the probabilities of the scenarios, requiring the values of the uncertain parameters to be set prior to solving DMP. The first step of the approach for the TSSP is to collect the data needed to construct the empirical cumulative distribution function (ECDF) for each uncertain parameter. Then, the ECDFs are approximated with closed form expressions and utilized in a formulation of DMP proposed in the paper. The approach for the MSSP extends it by incorporating the estimation of conditional moments, followed by time series forecasting to construct the ECDFs, and then solving an NLP of DMP on a node-by-node basis. A secondary approach for the MSSPs determines uncertain parameter values at the nodes all at once and only the scenario probabilities are estimated using an LP of the DMP.

Scenario reduction approaches have also been studied extensively within the literature. Pflug (2001)[102] presents an optimization approach for scenario reduction. The objective function of the approach is to minimize the Wasserstein distance between the reduced set and the original distribution. The approach is used to generate a scenario tree representation of the uncertain process within the problem. Included in the approach is the ability to incorporate constraints within the tree construction to satisfy relationships and correlations present in the uncertain process.

Dupačová et al.[103] and Heitsch and Römisch (2003)[104] presented two heuristic algorithms for scenario reduction on a discrete and finite distribution of scenarios that follow a scenario-tree-

like structure for convex SPs. The algorithms reduce the scenarios through the removal, backward heuristic, or addition, forward heuristic, and reweighting. They provide bounds for Monge-Kantotorovich mass transportation problem that provides the optimal weighting of a reduced scenario distribution to minimize the upper bound of the Fortet-Mourier metric. The algorithms presented in these works are the basis of the SCENRED tool found within the GAMS library. Further studies[105–107] were performed and used for the improvement and release of the SCENRED2 tool. Heitsch and Römisch (2007)[105] extended their original work from the upper bound to the direct usage of the Fortet-Mourier metric to improve the resulting quality of the reduced scenario set. The work reveals that using the Fortet-Mourier metric improves the reduced scenario set compared to the upper bound of the metric because the order of the metric is increased, implying a more stable solution for TSSPs. Heitsch et al.[106] investigate the stability of MSSPs with respect to filtration distances in addition to L_r-distances. The study found that only accounting for L_r-distances for stability based scenario reduction, though enough for TSSPs, are not sufficient for MSSPs, and the addition of filtration distance should be considered to attain stable reduced scenario sets for MSSPs. This finding paved the way for a later study[107], in which a scenario reduction algorithm is presented that accounts for both metrics within the optimization problem. The algorithm recursively removes a single node from the original scenario tree. Their findings demonstrated that the incorporation of the filtration distance into the problem provided a noticeable reduction in the final number of scenarios on the tree when compared to just using the L_r-distance.

Li and Floudas (2014)[108] present the optimal scenario reduction, OSCAR, framework. This framework is an optimization-based approach for determining a reduced scenario set. Like Dupačová et al.[103], the framework considers the Kantorovich distance between the original and reduced scenario set. In the objective function, however, OSCAR also considers the difference

between the model output for the reduced and full sets and the best and worst performing scenarios. The objective function is formatted so that each metric can be assigned a different weight to emphasize one metric over another. One drawback of this framework is the computational power required to calculate Kantrorovich distance. In a later study, Li and Floudas (2016)[109] extend their original framework to account for this drawback through a sequential scenario reduction approach that can quantify the quality of the reduced scenario set even when the original scenario set is too large to evaluate through factorial combination. A further extension by Li and Li[110] removed the necessity that the scenarios be generated from the factorial combination by incorporating clustering within the reduction algorithm.

Karuppiah et al.[111] propose a heuristic approach for scenario reduction using a MILP. The objective is to minimize the number of scenarios chosen for the reduced set. The decision variables are which of the original scenarios to include in the reduced set and the probability assigned to the included scenarios. To ensure a non-trivial solution, constraints were imposed to require that the probability of occurrence for each unique realization for each uncertain parameter remains the same between the original scenario set and the reduced scenario set.

It is unclear which of the various scenario set construction approaches would be the most effective for a problem. Hence, a number of studies evaluate different scenario construction approaches. Park et al.[112] aim to find the best scenario reduction method assuming the full scenario set was known. They tested four scenario construction methodologies, importance sampling, random sampling, the fast forward selection (FFS) algorithm utilized in the SCENRED tool[103–105,107], and their proposed stratified scenario sampling (SSS) approach. The methods were compared and evaluated on a two-stage stochastic power network investment problem. The study used three criteria, the expected cost of the naïve solution (ECNS), the solutions generated by the

methods, and a regret measure. It concluded that the FFS and SSS methods provide solutions that closely match the objective value of the true solution with a significant reduction in solution time, the decisions determined by FFS closely match those of the full scenario set, and finally, the proposed SSS method provides less accurate solution than FFS but a lower worst-case regret than FFS.

Li et al. (2020)[113] reviewed nine scenario construction methodologies used in the energy systems SP literature. They outlined the characteristics of the methods and assessed them based on the solution stability. They concluded that, in general, sampling-based methods were fast and simple, but the methods scaled and performed poorly when the problem uncertainty began to increase in dimensionality. Forecasting-based methods, methods where scenarios are generated from models trained using historical data, were found to accurately capture the trends in the data, such as correlations of variables and nonlinear relationships. However, as with many data-driven methods, forecasting-based scenario construction methods only perform as well as the available data. Lastly, the authors concluded that optimization-based methods provide a high degree of accuracy for the scenario space approximation for a wide range of problems, with the drawback that as the SP increases in complexity, the optimization problem to construct the scenarios also becomes more computationally demanding to solve.

Recently, Seljom et al.[114] evaluated six scenario construction methodologies. The methods included were 1) random Monte Carlo sampling, 2) iterative sampling, where a large number of scenario sets are constructed and the overall distance from the first four moments is used to determine the best set, 3) k-means clustering algorithm, 4) a constrained k-means algorithm that constrains the size of the clusters, 5) an optimization-based method for minimizing the distance of distribution moments and correlation, and 6) an optimization-based method for minimizing the

Wasserstein distance. To evaluate the quality and stability of the constructed scenario sets, the authors compare the methodologies based on the error in the distribution representation and the quality of the solution. The distribution error was measured in two ways. When the distribution error is measured based on moment matching, it was found that the iterative sampling scheme provided the lowest error. However, when the distribution error was evaluated using the Wasserstein distance, both k-means clustering algorithms performed the best. The evaluation of solution quality only compared the random sampling, iterative sampling, and the moment matching optimization methods due to temporal constraints. For the solution quality comparison, the authors noted that the optimization methodology performed the best providing the lowest and most stable optimality gap of the three methods. Additionally, they concluded that the iterative sampling approach yielded more stable solutions than random sampling. These reviews provided a good rule of thumb which scenario construction methodology performs well with respect to the metric deemed most important for the approximation of the scenario space. However, the only study to consider the impacts of rare events on the scenario construction methodology was Park et al.[112] through their evaluation of the regret-based metric.

Though much of the scenario construction literature has neglected the assessment of the impact of rare events within the scenario set, a handful of studies have investigated the incorporation of rare events within an SP or at least assessed the impact a given set scenarios had on decision variables. Ehrenstein et al.[115] propose a two-step scenario reduction algorithm for a multi-objective supply chain planning problem with extreme events to reduce the number of scenarios while including rare events. The algorithm utilizes the SAA approach but exploits the structure of the scenarios with rare events. The authors only use two scenarios instead of sampling n scenarios for the reduced scenario set, the worst case scenario, where all possible disruptions for

the supply chain happen at once, and the nominal case, where no disruptions occur. They then generate upper and lower bounds for the objective of the problem. The problem is then solved using an augmented epsilon constraint method to find the Pareto-optimal solution by using the two-scenario reduced set to determine the first-stage decisions and the full scenario set with the fixed first-stage decisions to obtain the second-stage decisions. The results revealed that the approach approximated the true trade-off solutions while reducing the solution time by around 99%.

Garcia-Herreros et al.[116] used a specialized solution approach using a strengthened Benders decomposition. First, the problem formulation was tightened using the convex hull of disjunctive constraints. The tightened problem was then solved using a multi-cut Benders decomposition to provide as much information as possible to the master problems from the subproblems. One way the authors strengthened the algorithm was to include information from the most probable scenario, where the supply chain is fully operational, in the master problem. The included constraints provided a better lower bound as the included information from the scenario had a sizeable impact on the overall solution. The second way the algorithm was strengthened was to ensure that a strong set of Benders cuts were included during each iteration. The strong cuts were generated by solving an LP that ensures the generation of Pareto-optimal dual multiplies for the sub-problems. The authors also included an approach to calculate the bounds on the penalties accrued for infeasible decisions due to using a reduced scenario set. The results showed that the solution approach reduced solution time by nearly 20 times that of solving the model directly, and the bounds and solutions produced by the reduced scenario set were very similar to that of the full scenario set problem. Both of these approaches, however, are problem specific and cannot be utilized for every problem.

Feng and Ryan[117] presented a heuristic scenario reduction methodology that extends the FFS algorithm[103–105,107] to incorporate the impact scenarios have on decision variable values, but not solely evaluating rare-event cases. The proposed methodology first requires optimal second-stage decisions to be determined for the full, or very large, scenario set to calculate the objective function as if each scenario were a deterministic problem given a set of feasible first-stage decisions. Next, the objective values of the set of scenarios are scaled to have similar magnitudes, and then a clustering method is applied to generate a specified number of clusters. Finally, the reduced scenario set is constructed using the FFS algorithm to select a single scenario from each cluster. The authors compared their approach to the original FFS algorithm on the stochastic reliability unit commitment problem. The results revealed that their proposed method constructed a scenario set with similar solutions to the ones obtained by the FFS. The solutions were also more reliable, i.e., they had lower shortages, and their method generated the scenario sets quicker than FFS.

2.3.2 Uncertainty Types within Stochastic Programming Framework

Uncertainty within an SP can be broken into two main classifications, exogenous and endogenous[118]. Exogenous uncertainties refer to the uncertainties the decision-maker has no control over, such as the weather. On the other hand, endogenous uncertainty refers to the uncertainties that a decision-maker can impact, such as production rate from an oil well or the failure rate of a supply networkwork linkage. Endogenous uncertainty has been further classified into type I and II based on how a decision impacts the uncertainty[119–121]. Type I endogenous uncertainty, also referred to as decision-dependent probability, occurs when the decisions impact the probability or the distribution of uncertainty. For example, if a planner invests in the strengthening of a network link the probability of failure will be reduced. In Type II endogenous

uncertainty, the decisions impact the timing of the realization of the uncertainty. As an example, the production rate of an oil well would vary based on when and what pump type a field developer installed a well. It should be noted that type I and II endogenous uncertainties can occur simultaneously, resulting in type III endogenous uncertainty[122] where decisions affect the realization timing and the distribution of the uncertainty. The CRCSP problem includes both exogenous and type I endogenous uncertainties.

A majority of the focus in the literature thus far has been on solving problems with exogenous and/or type II endogenous uncertainty. However, in recent years, there has been an increase in problems incorporating type I endogenous uncertainties within the SP literature, though it is relatively sparse compared to problems with other forms of uncertainty. Though theoretical models have been discussed prior[123,124], the first study to explicitly model this form of uncertainty is Ahmed[125]. The problem investigated was a p-choice location planning problem, where the probability of using a given location is based on the fraction of that location's utility compared to the system's overall utility. The author presented a solution approach to a class of one-stage SPs that used Luce's choice axiom for the probability measure for each scenario. The approach implemented a branch and bound algorithm where a reformulation of the 0-1 hyperbolic program at every node was performed and a genetic algorithm to generate the lower bounds.

One of the most common problem types incorporating type I endogenous uncertainty is disaster planning over a given network[126–134]. The first of the studies that investigated this problem was Peeta et al.[126], where the problem aimed to strengthen a highway network subject to linkage failures due to natural disasters. The probability of linkage failure due to the disaster is known and can be reduced by investing resources for a link. The problem was formulated as a TSSP, where the first-stage decisions were the resource allocation for the network strengthening, and the

second-stage problem minimizes the traversal costs from an origin to a destination of the network given the disaster realization. The authors approximated the objective function using a first-order Taylor series expansion to solve the problem. The reformulation led to a knapsack problem that could be solved directly, and the solution produced was guaranteed to be a local optimum. As a follow-up, Du and Peeta[127] extended this problem by incorporating the uncertainty of the disaster severity. A new solution approach was developed for this problem, called the Two-Stage Heuristic Algorithm. The approach split the TSSP into two separate problems and iteratively solved them until solutions to both problems converged. The algorithm could quickly converge to a feasible solution approximating the optimal solution for moderate-level disasters. However, the approach underestimated the solution for the high severity disasters and overestimated it for low severity disasters. Laumanns et al.[131] revisited Peeta et al.'s original study[126] and proposed a reformulation that could be directly solved without introducing non-linear terms. They proposed a technique labeled the distribution shaping technique for SPs, where binary decision variables impact the probability of an outcome. The technique derives a polyhedral characterization of the probability distribution that scales the probability measure up or down depending on a linear relationship between the binary variables. This reformulation and a proposed scenario bundling technique allowed the authors to solve the problem in a matter of seconds compared to the original solution times, which were on the order of minutes.

In the thesis of da Costa Flach[133], the motivation was to solve the resource allocation problem for strengthening a network to reduce the failure probability under disasters and the following disaster response transportation problem after the realization of the disaster. Much like Peeta et al.[126], this problem was formulated as a TSSP, with the first-stage decisions being the network strengthening decisions and the second-stage decisions the disaster response. The

uncertainty was the probability of a network link failing when a disaster occurs, where the probability of failure could be reduced by allocating resources to strengthen that link. The study investigated different problem reformulations, applied a cut generation algorithm, and developed a scenario generation schema to solve the problem. The reformulation technique to eliminate the effects of type I endogenous uncertainty relied on the decision variables being binary and the expression for the probability of failure for each link being linear. Da Costa Flach proposed using the property of logarithms and introducing a new variable to represent the natural logarithm of the scenario probability to linearize the probability measure. A set of piecewise linear terms was used to approximate the exponential function to convert the log-probability to an estimated scenario probability. The cut generation algorithm iteratively solved the linearized model adding more piecewise linear cuts to approximate the exponential function until the solution converged to an accepted error level. For scenario generation, da Costa Flach proposed the use of importance sampling to reduce the number of scenarios needed to obtain a stable estimate of distribution properties. The results revealed that, given the proposed approaches, models of different sizes might be solved directly to within 1% global optimality within a reasonable time. At the time, most studies only solved this class of problems through heuristic approaches to obtain good answers rather than optimal.

Medal et al.[128] presented a model for resource allocation for disaster protection of a resource network. The problem was once again presented as a TSSP, where the first-stage decisions were the protection resources allocated to the various facilities in the network, and the second-stage decisions correspond to the transportation of goods to meet demand post-disaster. The probability of capacity reduction realizations and the possible network configurations were impacted by the protection level given to each node and arc in the network as governed by the first

stage decisions. They reformulated the problem to eliminate the inherent nonlinearity within the objective functions of SPs with type I endogenous uncertainty with multiple decision variables in the first stage, allowing them to solve the problem to optimality with the L-shaped decomposition methodology[135]. A greedy algorithm was used to solve the model. Capitalizing on their model's properties, namely the submodularity of their model, the authors incrementally allocated protection based on the marginal benefit. The greedy algorithm solved the model substantially quicker than the decomposition method and was found to yield solutions that were at worst 5% off from the true solution.

Bhuiyan et al.[134] further investigated the problem in Medal et al.[128] by allowing the construction of the network itself. The problem, once again, was formulated as a TSSP. The first-stage decision variables were the resource allocation for protection and construction of links from facilities to demand locations. The second-stage decisions were the distribution decisions after a disaster. Due to the introduction of the network design in the first stage, the solution approach from the previous work[128] was not applicable. The model was linearized using a similar procedure[128]; however, extra linearization constraints were needed due to the transportation problem now being dependent on the first-stage decisions. The solution methodology used an accelerated Benders decomposition, where the acceleration was achieved by introducing multiple cuts per iteration. The additional cuts were analyzed to determine the combination of cuts that resulted in faster solution times. They concluded that multiple optimality cuts, trust region cut, Pareto-optimality cut, master problem node limit, and introducing a valid inequity formed the best combination.

Krasko and Rebennack[130] presented a model for allocating resources to mitigate and prevent post-wildfire destruction from debris flow and provide emergency response routing within

a network. First, the study presented an MINLP to minimize the expected damage due to debris flow, where the probability and the damage associated with the scenario are functions of the mitigation decision variables. This model is then combined with an emergency response problem to form a TSSP. The first stage problem allocates resources to reduce damage and the probability of scenarios where debris flow cause damage occur and also allocates emergency response vehicles across a network. The second stage problem is a multi-period vehicle routing problem to respond to the debris flow given the first stage decisions. The objective is to minimize the expected cost of the damages associated with the disaster, where the value of life is included to assign a monetary value to the emergency vehicle routing problem. The MINLP model was solved using the global solver BARON[136].

Escudero et al.[132] presented a three-stage SP for natural disaster mitigation with type-1 endogenous uncertainty and an algorithm to solve the problem. The objective was to minimize the costs associated with disaster preparation and response, including penalties due to shortage of demand and going over budget. The first-stage decisions include investments to obtain more accurate estimates of the scenario probabilities, thus introducing type-1 endogenous uncertainty, and the allocation and capacity modification of commodities within the supply chain. The second-stage decisions reinforce facilities to protect against disruption and include a second round of commodity allocation and capacity modifications. The third-stage decisions correspond to the disaster response of supplying commodities to the demand points in the network. The Cluster Dual Decent Algorithm (CDDA)was developed to solve the resulting problem, the algorithm relies on duality theory to improve incumbent solution. The CDDA starts first by solving for solutions of a relaxed version of the problem. From there two iteration schemes are used, and inner and an outer. The inner iteration seeks to identify local optimum through evaluating the impact to dual value of

the problem by selecting an additional facility. On the other hand, the outer iteration seeks the global solution by restructuring the solution generated by the inner iteration by adjusting capacity levels by solving the dual problem. The CDDA was shown to solve the problems at worst ten times faster than CPLEX while achieving around a 1% increase in the objective function value.

A second problem modeled with type I endogenous uncertainty in the literature regards maintenance planning and scheduling[137–139]. One of the first studies to do so was by Ekin[137], where the author presented a form of the planning and scheduling of an integrated maintenance and production planning problem. The problem was modeled as a TSSP, where the first-stage decisions involved system maintenance, and the second-stage decisions were production planning. The objective was to minimize the cost associated with maintenance and production. In this study, the percentage of the product that was on specification was modeled as a continuous uncertain parameter whose distribution parameters were dependent on the first-stage decisions. Ekin utilized an approach coined augmented probability simulation[140] to solve the problem due to the problem's objective function not having a closed-form expression. This study demonstrated that a continuous distribution, such as the normal distribution, can be included within an SP problem even when the uncertainty is decision-dependent, and it provided a solution approach.

Leo and Engell[138] presented a model that considers both production planning and plant maintenance, where the estimated remaining useful plant life depends on the plant operations. The probability of the remaining years of services was described through survival analysis. A degradation trajectory was incorporated into the survival analysis based on the operating conditions of the plant. If the plant operated at full capacity, the degradation rate was greater than if it operated at half capacity, causing the failure to occur faster. The problem was formulated as an MINLP and solved using three different approaches using BARON with a custom branching

scheme, a generalized Benders decomposition algorithm, and the Global Optimization Algorithm (GOP)[141]. The study concluded that the Benders decomposition was the most efficient solution approach when considering both global optimality and the solution time.

Yin et al.[139] presented a study that aims to determine the optimal maintenance schedule for a large scale wind farm that maximizes the revenue. The model is formulated as a TSSP where the first stage actions are the timing of the maintenance decisions and the second stage variables are the resulting power produced from the wind farm. In the problem, the power generated from a wind turbine depends on that turbine's availability and the surrounding turbines to consider wake effects. The probability of turbines being available depends on the first stage maintenance actions. Due to the interrelated complexities associated with the maintenance decisions, power generation, and the wake effects, along with the non-linearities and non-convexity of the problem, the authors employed a DOF, particle swarm optimization, to solve the problem. The solution revealed that by incorporating the wake effects seen in large scale wind farms, an increase in the expected net revenue can be observed and that the implementation of a maintenance strategy, regardless of the incorporation of the wake effects, results in a substantial increase in the net revenue.

Supply chain resilience and production planning is another problem with type I endogenous uncertainty[122,142–145]. Hu et al.[142] investigated system reliability and cost minimization for meeting power demands. The study considered the probability of a forced outage to be dependent on the operating conditions. A TSSP was developed to model the problem, where the operating conditions, units that are online, and the operating state were determined in the first stage, and the production rate of the units was determined in the second stage to meet demand. The scenario probability calculations were highly nonlinear due to the dependency on the operating conditions. To solve the problem, the authors presented a heuristic algorithm referred to as the adaptive

reliability improvement unit commitment (ARIUC) algorithm. In this algorithm, the probabilities were first assumed to be constant, and the problem was solved to obtain an initial solution. Then, the reliability of the system was evaluated, and the decision-dependent probabilities were reintroduced into the problem. At this point, the algorithm entered into a dual-loop to improve the solution if the system reliability did not match the termination criteria. In the loop, the algorithm determined and replaced units highly sensitive to generation capacity with units uncommitted from the previous solution. The problem was solved again for the second stage values, with the resulting solution evaluated for convergence of the system reliability. The authors concluded that the incorporation of the type I endogenous uncertainty provided a solution that reduced the overall cost and the loss of load probability associated with the system.

Another power generation expansion planning problem was presented by Zhan et al.[143], where the electricity price depended on investment decisions. The study modeled the problem as an MSSP where, at each stage, a decision may be made to invest in increasing power generation capacity. The objective was to maximize the net revenue of the power generation system. In the model, the outcomes at each stage could be either a high or a low electricity price, where the probabilities of the different outcomes were a function of the cumulative power generation capacity of the system. The problem was reformulated to eliminate the nonlinearity due to the decision-dependent probability. The reformulation represented the continuous 0 -1 variable of the probability of an outcome by approximating the value using a sum of binary variables and an error term. These new values were then directly linearized with the continuous variables of the objective function yielding a linear model. The authors solved the nonlinear MSSP using BARON and approximated linear MSSP using CPLEX. The results showed that the approximate linearization was solved much faster than the original nonlinear model, while the difference in the solutions was

negligible. Additionally, they found that incorporating type I endogenous uncertainty provided better solutions than the model with only exogenous uncertainty.

The study by Ma et al.[144] modeled a power grid resilience problem under uncertainty as a TSSP with type I endogenous uncertainty. The objective was to minimize the investment and expected downtime/repair costs given a grid-disrupting event. The first stage decisions consisted of line hardening, additions of distributed generation units, and adding line switches within the grid. The second stage decisions were operational and used to mitigate the weather disruption for the grid. The uncertainty included the grid load, wind-based disruption events, line repair time, and line damage status. The line damage status was impacted by the first-stage line hardening decisions, while the other sources were unaffected by the decisions. Two progressive hedging (PH) algorithms were proposed to solve the model. The first one solved each scenario problem independently and aggregated the individual solutions. The second PH algorithm performed a scenario bundling procedure to reduce the number of scenarios. The results revealed that PH could solve large-scale models of up to 123-bus distribution systems and yielded higher quality solutions with scenario bundling.

Tong et al.[145] studied a supply chain planning problem for the oil sector under product demand and yield uncertainties. The model was formulated as a TSSP where the first stage decisions defined switiching between the operational mode and the second stage decisions were the production considerations to meet demand, the problem also included CvaR constraints for meeting product demand. The study employed a Markov chain to describe the product yield fluctuations, where the fluctuation probabilities were tied to the operational switching decisions. The author presented an iterative solution heuristic based on the SAA approach to provide several solutions with varying degrees of risk aversion. The study found that by utilizing a Markov chain

representation of yield fluctuation, the model could reduce both costs and the risk of not meeting demand than using a uniform distribution.

In a study by Hellemo et al.[122], a general procedure was introduced to incorporate continuous distributions within a TSSP where the first stage decisions were the parameters of the endogenous uncertainty's probability distribution. The example problem used was a capacity expansion problem to meet power generation demand, where either the demand or unit cost was stochastic. The first stage decisions impacted the probability of the scenario occurrence and the capacity expansion decisions, while the second stage decisions defined the system production. The paper presented three approaches for modeling decision-dependent probabilities within SPs. The first was linear scaling of a uniform distribution where decisions impact the scaling factor. The second approach employed a convex combination of probabilities where decisions impacted the weight of each probability in a mixture model. The third approach used an approximation or the direct expression for a CDF where distribution parameters were decision variables. The authors investigated the impact the three approaches had on solution time and quality of BARON as the approaches provided the solver with specific problem information. They also employed a generalized Benders decomposition (GGBD) approach to solve the problem. The results revealed that providing insight about the model to the solver (BARON) improved its performance even with the complexities of type I endogenous uncertainty, demonstrating that problems with type I endogenous uncertainty could indeed be solved directly. The authors note that BARON was often much slower than the GGBD approach, however the results of the GGBD were not included in the study.

A few studies considered problems that do not fit into any of the problem groups discussed up to now[146–148]. One such problem was studied by Chen et al.[146], who combined a system

reliability problem with a process synthesis problem. The object was to minimize the total annualized cost while considering the possibility of system failure. The problem was formulated as a TSSP where the first stage decisions determined which process to use, the number of unit redundancies, and the sizing of the process equipment. The second stage decisions defined the operational decisions of the overall process. The probability of the system being down was modeled as a type I endogenous uncertainty, where the system reliability was dependent on the redundancies for each unit, resulting in a disjunctive and nonlinear optimization problem. The solution approach modified the existing Logic-based Outer Approximation (LOA)[149,150] algorithm to account for the mixed-integer decision variables. The authors compared how the system design and operation changed when the problem was formulated as deterministic process synthesis, stochastic process synthesis, and a stochastic process synthesis with reliability using a test case. The results showed that incorporating system reliability provided a system with the highest profit, demonstrating the value of incorporating the system reliability within the process synthesis problem.

Bhuiyan et al.[147] investigated the impact of incentive uncertainty regarding landowners following through on forest fire prevention and mitigation. The problem was formulated as a TSSP. In the first stage, a budget was allocated to provide monetary incentives to landowners for fuel reduction treatment on their land. The second stage of the problem represented the damage seen by a wildfire given the fuel reduction outcome. The budget allocation in the first stage impacted the probability that the landowners would implement a fuel reduction strategy, resulting in type I endogenous uncertainty in the model. The probability of a landowner implementing the fuel reduction procedure was estimated using logistic regression. Three different models were constructed, (1) a risk-based allocation model, where the budget allocation was decided based on

risk measures for each land owner, (1) a uniform allocation model, where a landowner was offered the same budget as all the other landowners, and (3) a hybrid allocation model, where a predetermined subset of landowners was offered a fixed amount while the remaining were offered an amount based on the risk. The authors reformulated the probability measures using the probability chain approach[128], and the second stage estimated damage was evaluated using a simulation model for all possible scenarios. The resulting model was solved using Gurobi. The results showed that the risk-based allocation model was the most effective approach for reducing wildfire damages.

Zhang et al.[148] presented a TSSP with type I endogenous uncertainty to determine the surveillance strategy for prostate cancer patients within different cohorts and ages at diagnosis. The objective was to maximize the QALY gained through surveillance from a set of a predetermined list of strategies, determining which strategy to assign to which cohort. In this problem, the probability of detecting a later stage of prostate cancer was a function of the first stage decision variables and the probability of not detecting cancer in the previous period. The model was then linearized and solved using Benders decomposition to achieve a solution within a two-hour wall time restriction. The solutions had a higher gain in QALY than the recommended one-size-fits-all strategy to monitor prostate cancer remission.

2.3.3 Cancer Care Related Applications

The application of stochastic programming can be found in a wide range of fields, spanning from investment planning to medical decision making. However, the literature on the use of stochastic programming for cancer-related problems is relatively sparse. In addition to Zhang et al.[148], Alvarado and Ntaimo[151] presented three stochastic programming formulations for scheduling chemotherapy appointments to minimize patient wait time and nurse overtime. The

decisions were the nurse schedules and times to schedule the patients. They considered the uncertainty of the appointment length, patient's level of care needed, or acuity level, and the total number of available nurses. They employed a mean-risk form of stochastic programming, which favors risk-averse decisions. The study showed that the SP formulation could generate better solutions than current deterministic scheduling algorithms.

There have been several applications of using an SP framework for cancer treatment[152,153,154]. Sir *et al.*[153] investigated the radioactive dosage from radiotherapy, maximizing the radiation delivered to the affected tissue while minimizing dosage to healthy surrounding tissue over a treatment timeline. The objective was achieved by determining the radiation dosage regime for each appointment. The uncertainties were the setup position, which quantifies the location of the tumor in relation to the peak intensity of the treatment beam, and the radiation dosage at each sub-volume within the tissue. The study showed that the resulting treatment plans were improvements over the typical non-adaptive treatment plans.

Aside from cancer treatment, Vieira et al.[155] applied an SP formulation to a radiotherapy staff allocation problem to maximize patient throughput. The decisions were where the staff was allocated to provide necessary services for radiotherapy. The uncertainties were the waiting time, processing time, and delay of the various operations. The study found that using their formulation, the throughput of patients could be increased by around 6 % for the case study. There are various other cases of applying SP for problems associated with treating or allocating treatment resources for cancer. However, to the best of our knowledge, there has been no SP formulation to determine optimal screening strategies for the early detection or prevention of cancer.

We hypothesize that SP would be an effective choice for modeling and solving the CRCSP. An SP framework can represent decision-dependent uncertainties through the use of scenario

generation. This characteristic of SPs is well suited for determining optimal CRC screening policies because much of the uncertainty within the CRC screening planning problem is decision-dependent. The compliance to the screening test is based on the test chosen[156]; the health state at which a screening test detects depends on the individual and that individual's age. Finally, the positive or negative results can depend on the test type used[19]. All of these points can be properly modeled within a SP framework.

CHAPTER 3

SIMULATION OPTIMIZATION APPROACH

This chapter introduces the simulation optimization framework developed to solve the CRC screening planning problem. The overall framework (given in Figure 3.1) can be separated into a CRC natural history model and a screening section, which are implemented in a CRC microsimulation (MSM), and an optimization component, which generates various screening strategies that are passed to the microsimulation. For each strategy, the microsimulation constructs a cohort of individuals and their life histories without screening using the CRC natural history model. Then, each individual is screened according to the strategy altering their life history. The microsimulation is repeated multiple times for statistical significance. At completion, the microsimulation computes the cohort's expected gain in life span and expected costs associated with screening and cancer, which are passed to the optimization component for evaluation. Details of the MSM are presented in Section 3.1.

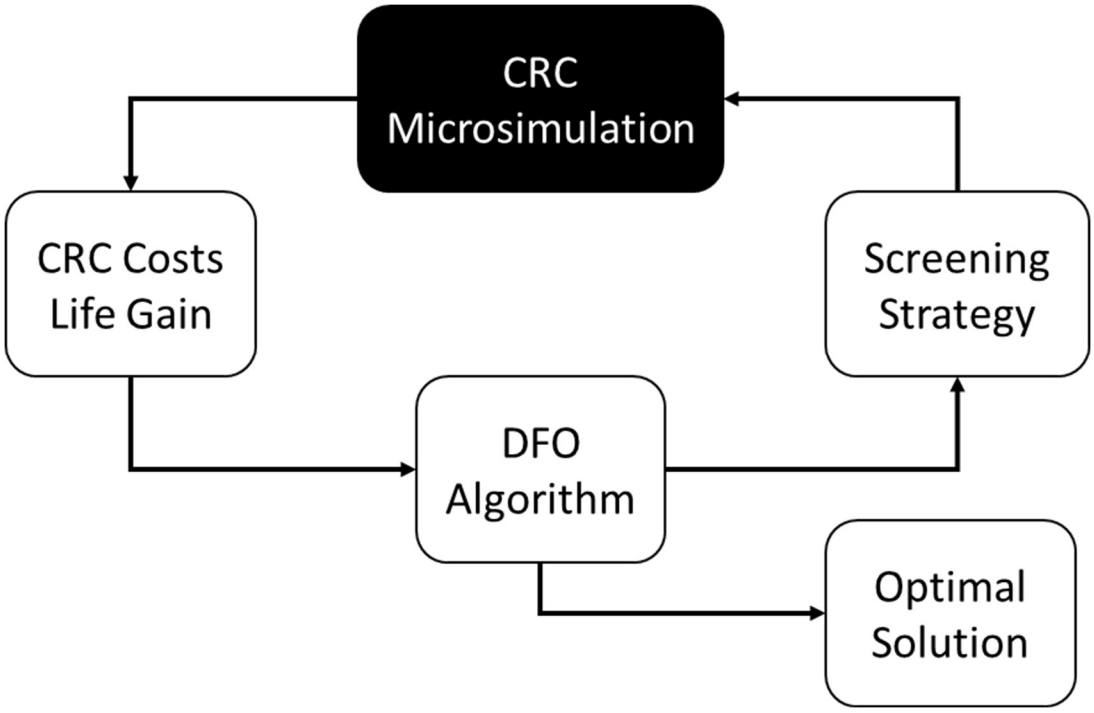

Figure 3.1 Simulation optimization framework for the CRC screening planning problem

The optimization component utilizes a DFO solver because the derivatives of the objective function, which is calculated using the CRC microsimulation is not available. As outlined in Section 2.2.1, there are numerous DFO solvers that can be employed to solve the CRC screening problem in a simulation-optimization framework. Therefore, the performance of these DFO solvers were compared using two test problems, optimization problems with similar characteristics to the CRC screening problem, and the results of these comparison study is discussed in Section 3.2. Next, a study was performed to investigate the impact that simulation assumptions and parameters

3.1 CRC Simulation

Of the microsimulation models within CISNET for colorectal cancer, Colorectal Cancer Simulated Population model for Incidence and Natural history (CRC-SPIN) is the most recently developed, in 2010, by the RAND Corporation[157]. This simulation model can be broken into two

different components, a natural history component and a screening component. The natural history component aims to recreate the progression of CRC within an individual free of any intervention and it is formulated as a continuous time model. The screening component allows for the disruption of the natural progression of CRC by early detection of CRC, or an adenoma, by the means of a screening test. Together the microsimulation creates two timelines per individual, relating to their life with and without screening, and compares the two to evaluate the effects screening had.

3.1.1 Natural History Creation Based on CRC-SPIN

The natural history portion begins by generating individual characteristics to define an individual. These defining characteristics are the individual's sex, age at the start of the simulation, age of death from sources outside of CRC, and CRC related risk parameters. Both sex and initial age at the beginning of the simulation are based around the desired population to model, with the percentages of the population described by age and sex combined to create an empirical cumulative distribution. Where random numbers are generated from a Uniform(0,1) distribution, the value corresponds to the resulting age and sex from the empirical distribution of the population. An individual's age of death from causes outside of CRC are generated using the U.S. life tables as a source for survival rate information based around the individual's initial age. This is determined by generating a value from a Uniform(0,1) distribution and incrementally stepping forward year by year until the survival rate is no longer greater than the generated value. Finally, CRC related risk factors are assigned through simulation specific parameter values, these parameters are denoted by $\alpha_{0i}, \alpha_1,$ and α_{2k}, where $k \in \{1,2,3,4\}$ signifying different age bins. Of these values, only one is individual specific, α_{0i}, which represents an individual's baseline risk for developing an adenoma. This parameter is generated through the use of two more simulation specific

parameters, α_0 and σ_0, as a randomly assigned value where $\alpha_{0i} \sim Normal(\alpha_0, \sigma_0^2)$. Values for $\alpha_1$and α_{2k} correspond to sex based risk and age based risk respectively.

Once an individual has their defining characteristics specified by assigning appropriate parameter values, the next step of the natural history model is to generate the ages at which the individual develops adenomas, if any. CRC-SPIN treats the occurrence of adenomas within an individual's life as a Nonhomogeneous Poisson Process, with the instantaneous risk of developing an adenoma, $\psi_i(t)$, modeled as a function of time as defined in Eq. **3.1**. Where $\boldsymbol{\delta}(\dots)$ is an indicator function equal to 1 when the expression given is true and 0 otherwise, $sex_{female} = 1$ or $sex_{male} = -1$, and $\boldsymbol{A} = \{20, 50, 60, 70, 120\}$ marks the ages that separate the four age bins.

$$\boldsymbol{\psi_i(t) = exp\left(\alpha_{0i} + \alpha_1 sex_i + \sum_{k=1}^{4} \delta(A_k < age_i(t) \le A_{k+1}) \left\{ age_i(t)\alpha_{2k} + \sum_{j=2}^{k} A_j(\alpha_{2\,j-1} - \alpha_{2j}) \right\}\right)} \quad 3.1$$

Using the expression for instantaneous risk, the total number of adenomas developed, *N*, can be treated as a homogenous Poisson Process across an interval (0,T] by means of integration, or $N \sim Poisson\left(\int_0^T \psi(u)du = \Psi(T)\right)$, where the expression for $\Psi(T)$ can be found in Eq. **3.2**

$$\boldsymbol{\Psi_i(t) = e^{\alpha_{0i} + \alpha_1 sex_i} \sum_{k=1}^{4} \left(\delta(age_i(t) > A_k) \left(\frac{e^{\alpha_{2k} \min(A_{k+1}, age_i(t))} - e^{\alpha_{2k} A_k}}{\alpha_{2k}} \right) \exp\left(\sum_{j=2}^{k} A_j(\alpha_{2\,j-1} - \alpha_{2j}) \right) \right)} \quad 3.2$$

Equation **3.3** describes the distribution related to an event occurring some ΔT amount of time from an arbitrary time point, t.

$$F(\Delta T) = 1 - \exp\left(-\left(\Psi(t + \Delta T) - \Psi(t)\right)\right) \quad 3.3$$

Adenoma occurrence dates for the individual, as well as total number of adenomas developed within a lifetime is simulated following the algorithm given in Figure 3.2.

START

1. SET $t = 20$
2. GENERATE $u \sim Uniform(0,1)$
3. LET $\Psi(t + \Delta t) = \Psi(t) - \ln(1 - u)$
4. SOLVE FOR $t + \Delta t = \Psi^{-1}(\Psi(t) - \ln(1 - u))$
5. IF $t + \Delta t > 120$ END
6. SET $t = t + \Delta T$
7. STORE t
8. GO TO STEP 2

Figure 3.2 Algorithm for generating adenomas within an individual

CRC-SPIN assumes that each adenoma is independent of others. As such, each adenoma has its own identifying characteristics. The first of the characteristics is the location of the adenoma within an individual. The location is assigned randomly according the distribution of adenomas found in Pickhardt *et al.*[53], where 8 %, 23 %, 24 %, 12 %, 24%, and 9 % of the adenomas occur in the cecum, ascending colon, transverse colon, descending colon, sigmoid colon, and rectum, respectively. The growth of an adenoma is set to follow the Janoschek model describing the diameter of adenoma j for individual i over time, d_{ij}, as given in Eq. **3.4**, where d_∞ is the maximum adenoma size, set to 50 mm, $d_0$is the minimum adenoma size, set to 1 mm, and $\lambda_{ij}$is individual i's jth adenoma's specific growth rate.

$$d_{ij}(t) = d_\infty - (d_\infty - d_0)e^{-\lambda_{ij}t} \qquad 3.4$$

The adenoma specific growth rate is determined by generating the time it takes for that adenoma to reach a diameter of 10mm, t_{10mm}, and back calculating the growth rate from the

Janoschek's growth model. CRC-SPIN assumes t_{10mm} is distributed by the cumulative distribution function found in Eq. 3.5, where $\beta_{1,l}$ and $\beta_{2,l}$ are simulation parameters where l is the adenoma's location, colon or rectum.

$$F(t_{10mm}) = exp\left(-\left(\frac{t_{10mm}}{\beta_1}\right)^{-\beta_2}\right) \quad 3.5$$

In the simulation, each adenoma is allowed to progress to cancer according to a log-normal probability distribution based around the size of the adenoma, s, and age at which the adenoma initiated, a. The cumulative distribution is shown in detail in Eq. 3.6, with γ_1 and γ_2 being simulation parameters that are categorized by the adenoma's location and the individual's age and $\Phi(\dots)$ being the standard normal cumulative distribution.

$$F(s,a) = \Phi\left(\frac{\ln(\gamma_1 s) + \gamma_2(a-50)}{0.5}\right) \quad 3.6$$

Since the age at which the adenoma is initiated is already defined, the size at which the adenoma becomes cancer is generated by Eq. **3.7**.

$$s_{preclin} = \frac{1}{\gamma_1} e^{0.5n - \gamma_2(a-50)} \quad 3.7$$

Where $n \sim Normal(0,1)$. With the size at which the adenoma becomes cancerous, an age is then back calculated using Eq. **3.4** for when the individual develops cancer from that specific adenoma. CRC-SPIN is then designed to assign a Sojourn time for that case of undetected or preclinical cancer. The Sojourn time, ST, is the interval between the cancer initiation and clinical diagnosis of that cancer, and is generated from a log-normal distribution with parameters ξ and ν, or $ST \sim LogNormal(\xi, \nu)$. The parameters ν and ξ are found using Eqns. **3.8** and **3.9** respectively, where μ and τ are simulation parameters that are categorized by the location of the adenoma within the individual.

$$\nu = \sqrt{ln(\tau^2 + 1)} \tag{3.8}$$

$$\xi = \ln(\mu) - \frac{1}{2}\ln(\tau^2 + 1) \tag{3.9}$$

If the age at which the adenoma becomes cancerous occurs before the individual dies of causes outside of CRC, the simulation treats that adenoma as a cancer from that point onward. In doing so, the size of the adenoma/carcinoma is no longer based around the Janoschek model. Instead, a size of the cancer at clinical detection is generated from a size distribution that is estimated from data obtained through the Surveillance, Epidemiology, and End Results (SEER) database. This data used are the diagnosed cases of colorectal cancer from 1975-1979[158], where the stage of CRC and the size of the tumor were recorded. The distribution of the data was estimated using the fit function of the python package scipy in Python 3.5. Due to the skewed nature of the data, the distributions considered for the estimated distribution were: truncated normal, generalized beta, log-normal, and gamma, using a chi-square test as the determination of best fit. The histogram of the data and the best estimated distribution is shown in Figure 3.3. It was determined that a generalized beta distribution, with parameters $\alpha = 4.17$, $\alpha = 6.98$, $location = -0.04$, and $scale = 1.34$, was the best fit, with a chi-square test p-value of 0.9999. With the size of the cancer at detection determined, an exponential growth model for the cancer is fit using the Sojourn time, ST, size of the cancer at clinical detection, s_{clin}, and the assumption that a detectible cancer begins at the size of 0.5mm. The cancer's growth rate, r_{cancer}, is calculated as shown in Eq. 3.10.

$$r_{cancer} = \left(\frac{s_{clin}}{0.5}\right)^{\frac{1}{ST}} \tag{3.10}$$

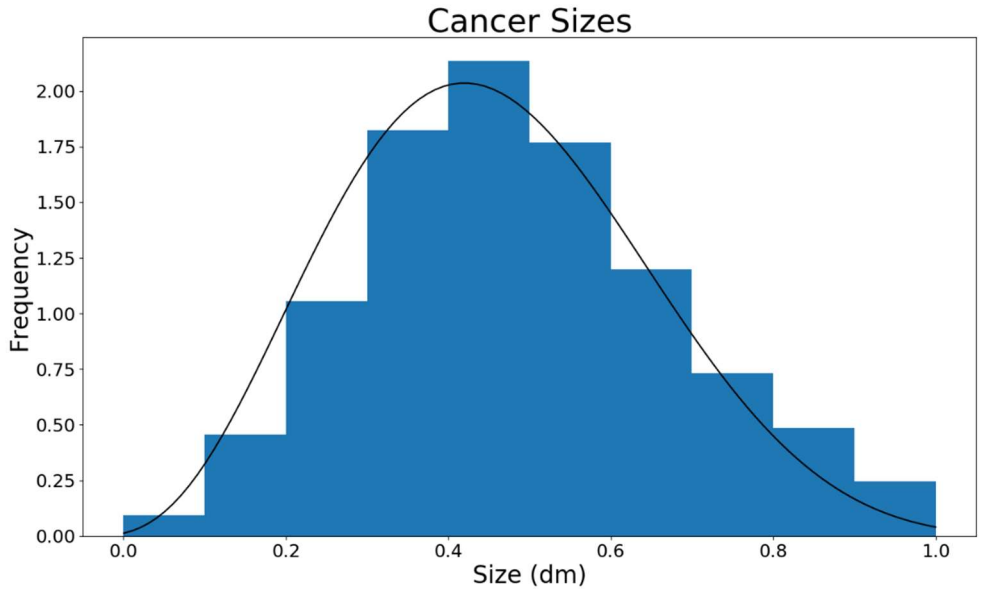

Figure 3.3. Distribution fitting for 1975-1979 CRC SEER data

Stage of cancer is assigned through a multinomial logistic regression (MLR) model, as was used by Rutter et al.[8]. This model uses the predictor variable as the size of the cancer at detection and the dependent, categorical variable as the stage of cancer, Stage I-IV. The MLR model uses a linear predictor function and takes the form shown in Eqns. 3.11-12, with $\beta_{n,i}$ being a regression coefficient corresponding to the nth term of the predictor function for the i-th category.

$$\Pr(stage = i \in \{1,2,3\} | size = s_{CRC}) = \frac{\exp(\beta_{0,i} + \beta_{1,i} s_{CRC})}{1 + \sum_{k=1}^{3} \exp(\beta_{0,k} + \beta_{1,k} s_{CRC})} \quad 3.11$$

$$\Pr(stage = 4 | size = s_{CRC}) = 1 - \sum_{k=1}^{3} \Pr(stage = k | size = s_{CRC}) \quad 3.12$$

The values for each regression coefficient was determined through the use of Maximum Likelihood Estimation (MLE)[159] utilizing CRC specific size and stage data from the SEER database[158]. Figure 3.4 displays the resulting fit of the MLR model as along with the data used to develop the model. The resulting fit shows a relatively poor fit at the lower range of sizes, however

after a size of about 20 mm we see the MLR model performing very well in representing the rest of the data.

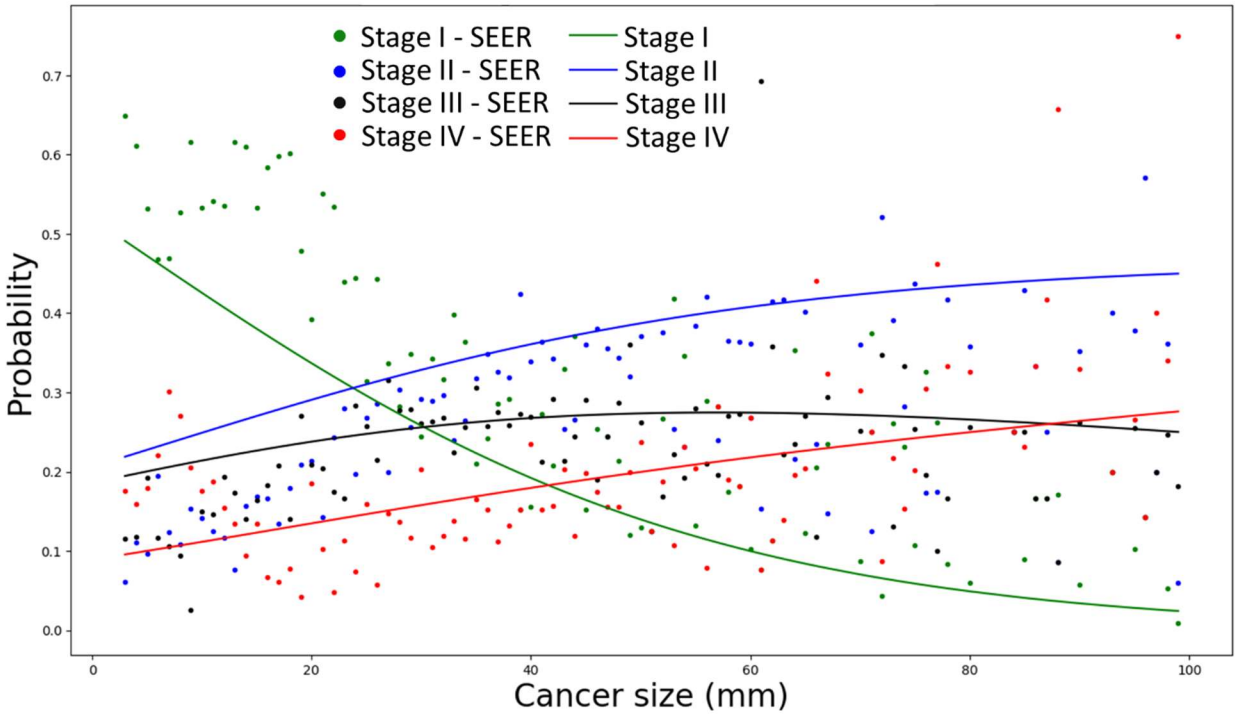

Figure 3.4 Multinomial logistic regression fit of CRC stage to size of the cancer

The final portion of CRC-SPIN's natural history section is the determination of death due to CRC. The model used to represent the hazard from cancer was the Cox's proportional hazard model[160] is found in Eq. 3.13.

$$\mathrm{H}(t) = \exp(\boldsymbol{X\beta} + \gamma_t) \qquad 3.13$$

Where $\boldsymbol{X}$ denotes a vector of covariates of the individual, $\boldsymbol{\beta}$ represents the hazard coefficient vector for the associated covariates of the individual, and γ_t represents the temporal hazard vector associated with time interval $(t-1, t]$, with t being the time since diagnosis. The covariates that were considered were those based on CRC-SPIN's initial model, age at diagnosis and sex of the individual. In addition to the covariates, a total of eight hazard models are developed based on whether the cancer is located in colon or rectum and the stage at which the cancer is

detected. The parameters of the model were estimated through the use of the NIH's CANSURV[161] software, with the data used for regression being CRC survival data from 1975 to 1979 within the SEER database[158]. With the hazard model constructed, the expression for the probability of survival up to interval $(t-1, t]$, found in Eq. 3.14, is used to generate the length of time lived after CRC diagnosis.

$$S(t) = \prod_{i=1}^{t} \exp\big(-H(i)\big) = \prod_{i=1}^{t} \exp(-\exp(\boldsymbol{X\beta} + \gamma_i)) \qquad 3.14$$

After the age of death due to CRC has been calculated for each of the adenomas developed by the individual, the simulation re-evaluates the size at detection for each of the carcinomas present. It determines the age at which the individual has the first clinically detected case of CRC and determines the size of all other preclinical cancers present at that age, evaluating the stage associated with the cancer based on the reduced size. The newly evaluated stage for each instance of CRC is then bounded by the stage that instance of CRC would be, were it detected from symptomatic means. This is done to ensure the probability of the new stage being more advanced than the stage if it were detected through symptoms. For the evaluation of assignment of age of death due to CRC, the age is determined by the instance of CRC that is the most advanced stage at the time of detection, shown in Figure 3.5. If there are more than one instances of CRC at the most advanced stage, the earliest age of death from CRC is taken. Finally, the individual's lifespan is updated, choosing the earlier of the age of death due to sources outside of CRC and age of death due to CRC.

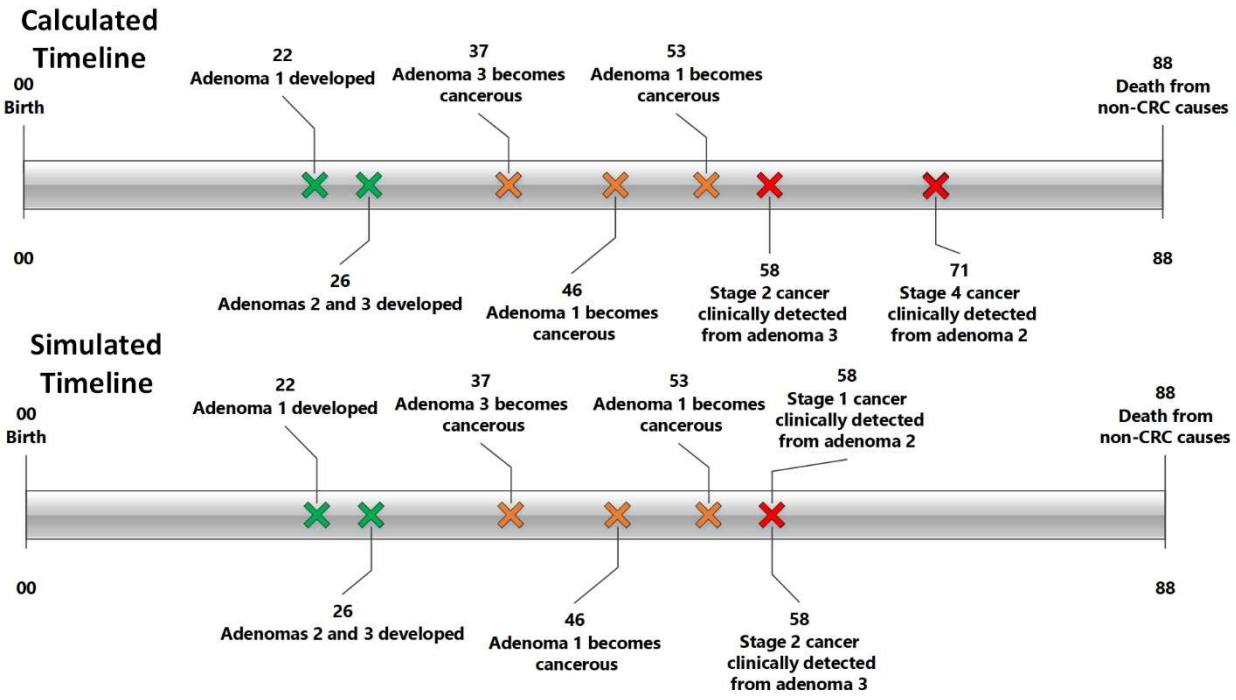

Figure 3.5 Example timelines demonstrating the model's determination of CRC stage at detection.

A total of 23 input parameters are used in the CRC-SPIN model. These parameters' distributions were determined through a Markov Chain Monte Carlo calibration technique[52]. In their work, Rutter and Savarino reported the prior distributions used for each parameter, the mean, the 2.5 % quantile, and the 97.5 % quantile of the posterior distribution, and the overlap of the posterior distribution with the prior distribution. As a way of estimating these posterior distributions, a wide array of probability distributions were evaluated using those four reported values. The parameters used to describe the trial distributions were evaluated by minimizing the sum of the square errors from the reported values. The distributions were ranked based on the sum of the square error around the reported values, and the distribution with the lowest error was used for each parameter. These parameters are listed with their associated use within the model in Table 3.1.

Table 3.1 CRC-SPIN's parameter list

Parameters	Associated model use
$\alpha_0, \sigma_0, \alpha_1, \alpha_{2,1}, \alpha_{2,2}, \alpha_{2,3}, \alpha_{2,4}$	Adenoma risk
$\beta_{1,c}, \beta_{2,c}, \beta_{1,r}, \beta_{2,r}$	Adenoma growth*
$\gamma_{1,c,m}, \gamma_{2,c,m}, \gamma_{1,c,f}, \gamma_{2,c,f}$ $\gamma_{1,r,m}, \gamma_{2,r,m}, \gamma_{1,r,f}, \gamma_{2,r,f}$	Transition from adenoma to carcinoma**
$\mu_c, \tau_c, \mu_r, \tau_r$	Sojourn time*

* Location specific: r = rectum, c=colon
** Location and sex specific: r = rectum, c=colon, m=male, f= female

With all of the individual sub-components of the natural history combined, the microsimulation generates an individual's lifetime, identifying all major events pertaining to CRC within the lifetime. A chart summarizing the flow of the natural history section for a single individual is given in Figure 3.6. This process can then be repeated a number of times to generate an entire cohort of people. Individual and population level trends from this unscreened cohort can then be calculated, e.g. annual CRC incidence rate, annual CRC death rate, number of adenomas in the population, or average number of adenomas developed by an individual.

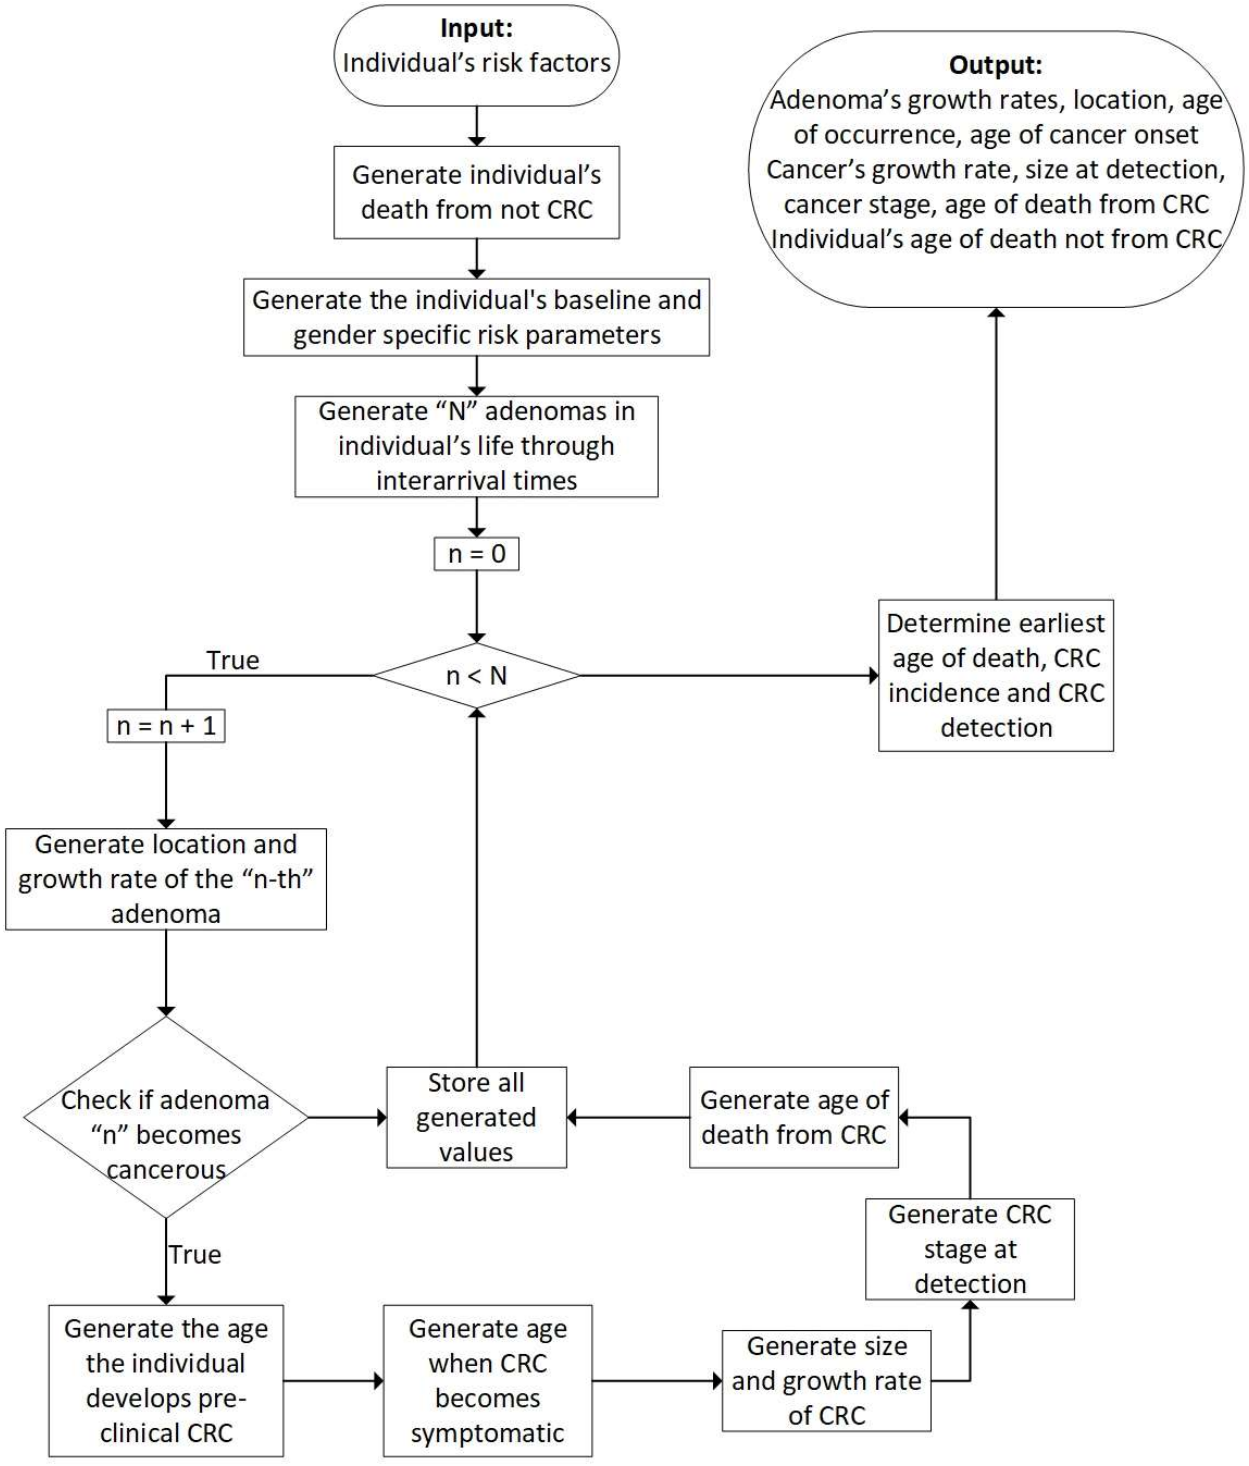

Figure 3.6 CRC microsimulation natural history component flowchart for a single individual

3.1.2 Verification of Implemented CRC Microsimulation

The first verification performed was to confirm the created CRC microsimulation the properly simulated adenomas within the individuals. Two outputs were investigated: 1) the rate of occurrence of adenomas within a population as a function of age, 2) the average number of adenomas within a 65 year old individual, given the individual had at least one adenoma, and the number of adenomas and their locations per 1,000 individuals generated.

To investigate the rate of occurrence of adenomas, a cohort of 15 million males were simulated, using the same individual base-line risk, $\alpha_{0,i}$, for each individual, storing the ages at which an adenoma was developed. A normalized histogram of the ages was then constructed, a plot of the instantaneous risk of developing an adenoma, Eq. **3.1**, was then overlaid on the histogram, seen in Figure 3.7. The generated distribution of ages, very closely fit the instantaneous risk function (solid orange line), confirming the implemented algorithm (Figure 3.2) provided a similar age distribution to that defined by Rutter and Savarino[8].

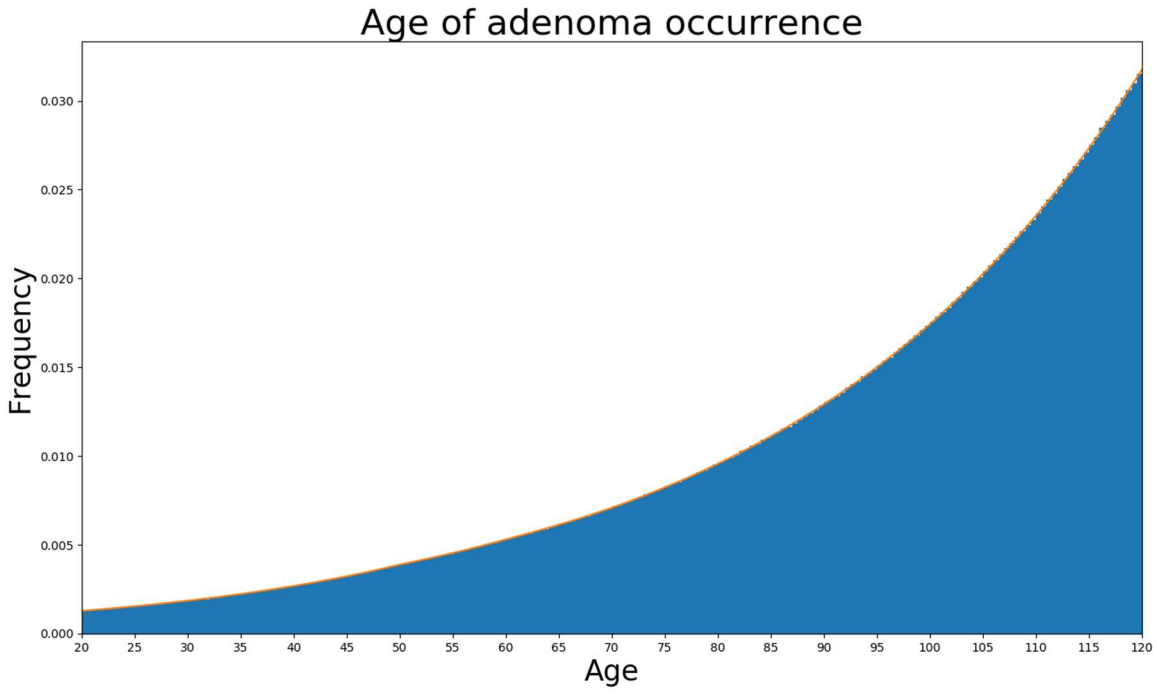

Figure 3.7 Adenoma occurrence verification

With the age distribution verified, we then had to verify the number of adenomas generated was approximately correct. This was verified by values presented in Table 3.2. This table investigates the total number of adenomas by size and location within a cohort of 1,000, 65-year-old individuals, showing a general under estimation of the number of adenomas generated. In addition to those adenoma totals, the average number of adenomas within a 65-year-old, given that individual has at least one adenoma, was calculated. It was reported[162] for CRC-SPIN for this value to be 1.8 adenomas, the value that was calculated from the CRC-SPIN recreation was $1.72\ (\pm 0.04)$. In general, the CRC model underestimates the total number of adenomas while accurately representing the rate of adenoma occurrence compared with CRC-SPIN. However, within the verification of CRC-SPIN[8], it was noted that CRC-SPIN on average overestimated the number of adenomas present. Thus the slight underestimation of the total number of adenomas developed is an acceptable deviation from CRC-SPIN, overall showing the created CRC microsimulation having a comparable simulation of the adenoma occurrence and growth to that of CRC-SPIN.

Table 3.2 Total adenomas by size and location

Location	Size	Reported[66]	Recreation
Rectum	1 – 5 mm	14	11.3 (± 0.6)
	6 – 9 mm	9	6.8 (±0.3)
	≥ 10 mm	27	18.8 (± 0.8)
Colon	1 – 5 mm	315	287.1 (±13.0)
	6 – 9 mm	112	102.0 (±4.5)
	≥ 10 mm	68	73.2 (±3.1)
Total	1 – 5 mm	329	298.3 (±13.5)
	6 – 9 mm	121	108.8 (±4.8)
	≥ 10 mm	88	92.1 (± 3.8)
	Total	538	499.2 (±21.6)

The second set of tests were performed to verify 1) the transition of adenomas to cancer and 2) cancer stage and location distribution at diagnosis. The sub-component of the transition probability of an adenoma to transition into CRC was verified through the simulation of 15 million

individuals, collecting the sizes at which adenomas transitioned to a cancerous state. Since the expression for transition probability, Eq. 3.6, is defined by location of the adenoma, sex of the individual, and age at which the adenoma developed, the stored transition sizes were broken into 16 sub-groups around those definitions. A histogram was then generated for each of the sub groups overlaying the corresponding cumulative distribution function, normalized to the maximum possible size of an adenoma. An example chart is displayed in Figure 3.8, showing the cumulative transition probability of an adenoma that developed in the colon of a female at age 45 as a function of size, similar plots of various ages for males and females for both colon and rectum can be found in Appendix A. The fit of this plot suggests our transition size is a valid representation of that found in CRC-SPIN.

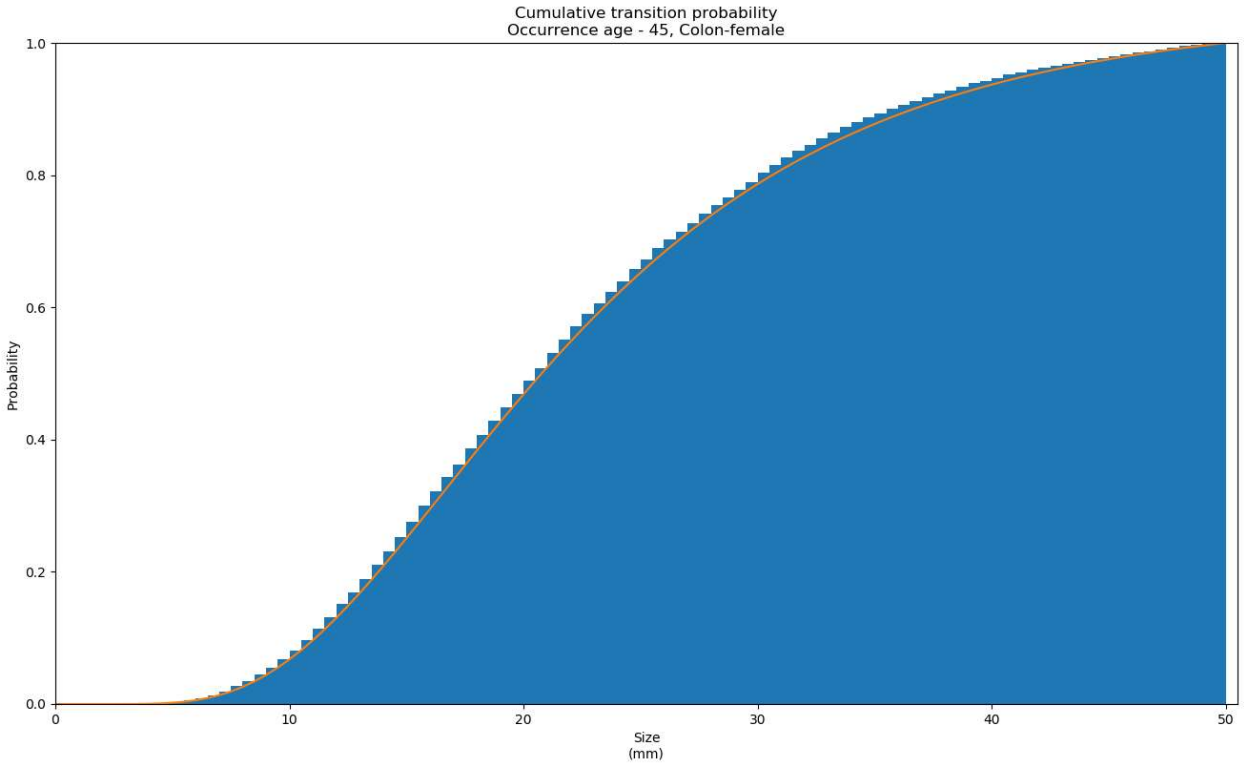

Figure 3.8 Cumulative transition probability of colon cancer for a 45 year-old female

To verify the distributions of clinically detected CRC, the same simulated 15 million people were used. When an individual was diagnosed with CRC, the location and the stage were stored. It has been reported[5] that approximately 30 % of all cases of CRC come from rectal cancer, the proportion of rectal cancer cases found from the created model was $29\ \%\ (\pm\ 0.001)$, showing our model provides a good representation of location distribution of CRC cases. The resulting distribution of the CRC stages at clinical detection for the created model and reported values of CRC-SPIN found in Table 3.3. The stage distribution of the created CRC model match that of CRC-SPIN's reported values, with slight deviation in Stage II and Stage III. However, the reported values of CRC-SPIN fall within the created model's 95% confidence interval. Both tests confirmed the CRC model properly represented the location and stage distribution of clinical cases of CRC.

Table 3.3 CRC stage breakdown

Cancer Stage	CRC-SPIN reported values[2]	Created model values
Stage I	18	18 (± .1)
Stage II	36	37 (± .1)
Stage III	27	26 (± .1)
Stage IV	19	19 (± .1)

The final verification tests involved the interactions of each of the subcomponents of the natural history model, investigating 1) average dwell times, or time needed to become cancerous, of the adenomas, 2) CRC incidence of a cohort, and 3) adenoma prevalence and CRC incidence as a function of age. All of these were tested by simulating 1,000 cohorts of 1,000,000, 55-year old, individuals, with each cohort having a different input parameter set generated from our fit simulation parameter distributions.

The dwell times investigated were the time it took for an adenoma to reach a cancerous state after it initially developed, the time it took for a cancerous adenoma to be clinically detected, also known as the Sojourn time, and the overall time for a cancerous adenoma to be clinically

detected after the adenoma develops. This calculation only considered the dwell times for clinically detected cases of CRC across all iterations. From the values collected the mean, median, and the interquartile range were calculated and compared to published[162] values for the CRC-SPIN model, presented in Table 3.4. The close match for our adenoma to a cancerous state dwell time is very promising, as it has been noted[162] that this dwell time has a very large impact on the screening efficacy. Matching these dwell times help confirm our model's representation of the progression of CRC in an individual. It should be noted that the mean and the median value for our recreated sojourn time are swapped from the published results. However, as the sojourn time is log-normally distributed, the median value should always be lower in value than that of the mean value, and it is believed these values were unintentionally swapped when reported.

Table 3.4 Dwell time information

Dwell time	CRC-SPIN values[162]	Recreated values
Adenoma to cancerous state		
Mean	24.2	25.01 +- 0.003
Median	23.0	23.38
Interquartile range	16 – 31	15.50 – 32.92
Sojourn time		
Mean	1.6	1.97 +- 0.0004
Median	2.0	1.59
Interquartile range	1 – 2	0.94 – 2.51
Overall		
Mean	25.8	26.98 +-.003
Median	24.0	25.37
Interquartile range	17 – 33	17.52 – 34.88

The cumulative CRC incidence investigates the fraction of the cohort that has been clinically diagnosed with CRC. To calculate the cumulative CRC incidence, the simulated cohorts were separated into two sub groups, one where no adenomas are present at age 55 and the other where an adenoma or preclinical cancer is present at age 55. Incidence was tracked for ages 55 through 85 for each sub-group as well as the number of individuals alive in each sub-group. For

this test, the risk of death from CRC was removed in accordance to the reported results[162] for CRC-SPIN. The generated cumulative incidence values were then plot against CRC-SPIN values, this comparison can be seen in Figure 3.9. The plot shows a good representation of the progression of CRC in the population without adenomas or preclinical CRC at age 55, but a moderate deviation for the representation of CRC progression in the population that had CRC precursors at or before age 55 is seen starting at around age 65. The general trend is still seen for both sub-population with respect to the cumulative CRC incidence, however the deviation for the population that had adenomas/preclinical CRC prompted the further investigation of the adenoma prevalence and clinical incidence rate.

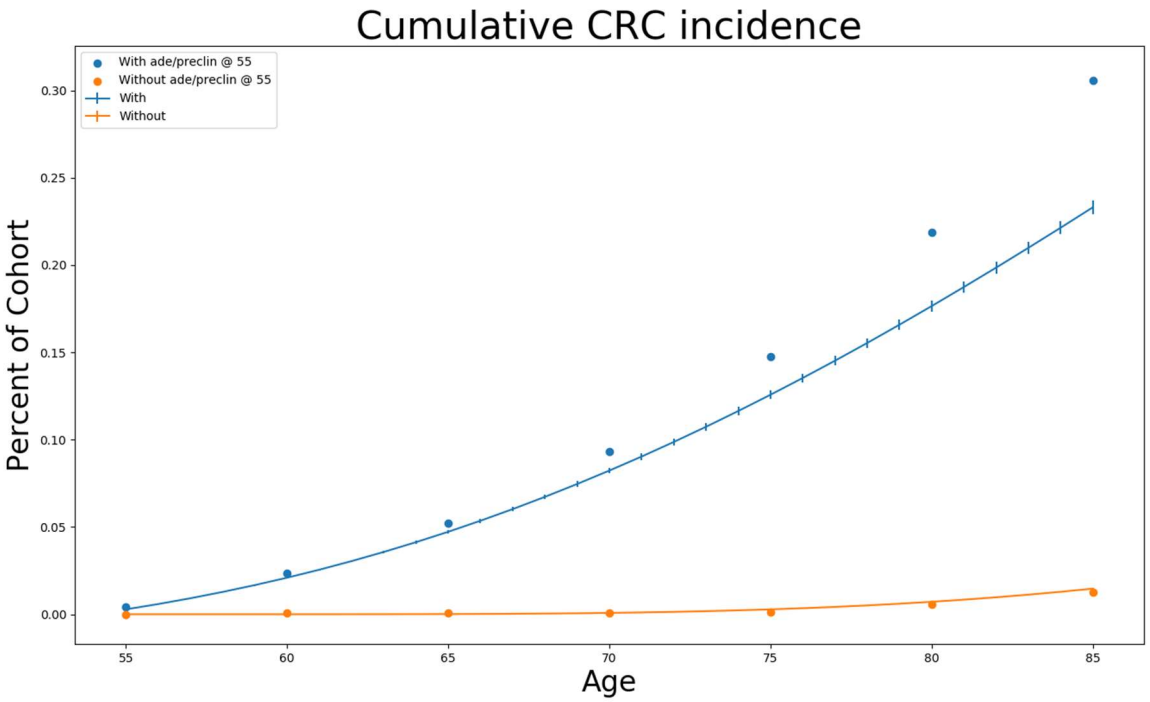

Figure 3.9 Cumulative CRC incidence[162]

The adenoma prevalence of a cohort measures the fraction of the population that has at least one adenoma developed at a given age, and clinical incidence rate of CRC denotes the total new cases of CRC diagnosed within the year. The number of individuals that had a clinically detected cancer was tallied based on age, and were then normalized to represent the rates per

100,000 individuals. The generated values along with their 95 % confidence interval were then plotted next to reported values[162] for CRC-SPIN, the plots can be found in Figure 3.10 and Figure 3.10 for adenoma prevalence and CRC incidence respectively. From the adenoma prevalence plot we notice a good match with ages less than 55 and greater than 80, then the range in-between, which was to be expected from the results presented in Table 3.2. It is also noted, by Figure 3.10, that the clinical incidence rate shows a similar deviation as that found in Figure 3.10. This deviation is likely a result from the under-estimation of adenomas presented in Table 3.2. The model over estimates the adenoma prevalence for the population ranging from age 55 to 80, but the total adenomas found at age 65 is lower than expected. This could easily contribute to the decrease seen in CRC incidence starting at age 65. Without a lower number of adenomas, that reduces the risk of an individual developing CRC. Although the model under-estimates CRC incidence in the later stage, while in general over-estimating the adenoma prevalence, the model was deemed acceptable deviation from reported values for the sake of our use within an optimization framework. This acceptable deviation was determined as the dwell times show that we still more or less capture that CRC progression within an individual's life time, meaning we should see reasonable results while applying our methodology for optimal screening determination.

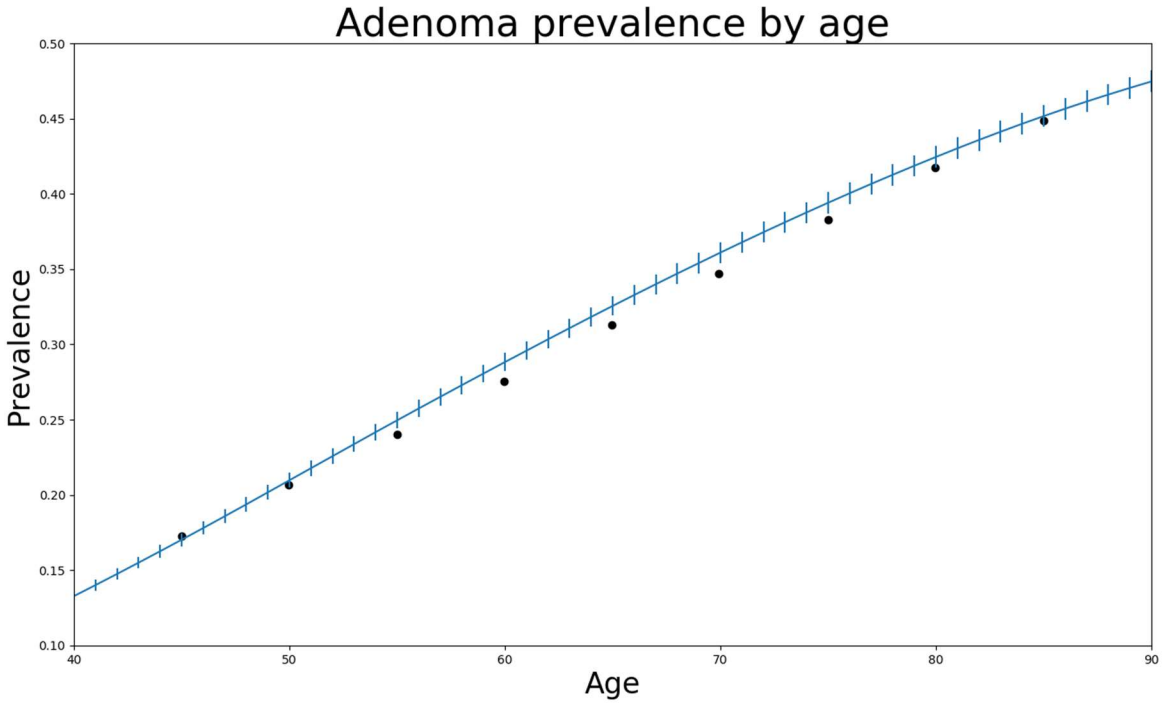

Figure 3.10 Adenoma prevalence as a function of age[162]

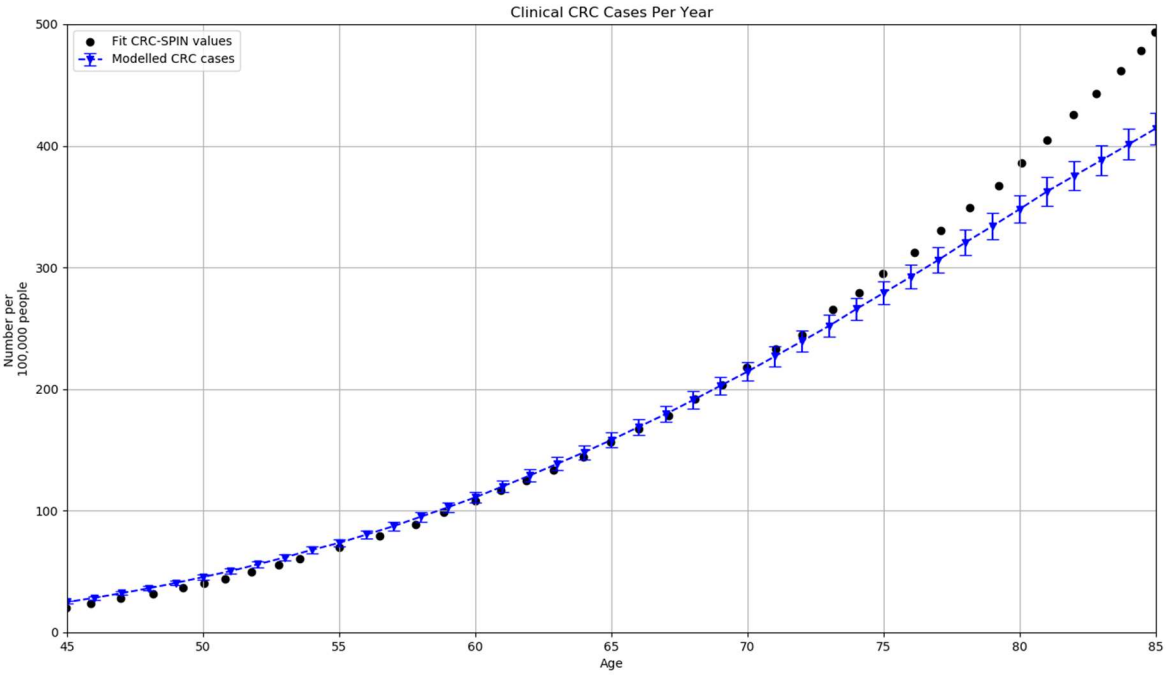

Figure 3.11 CRC incidence as a function of age[162]

3.1.3 Screening Component Implementation

The screening component of the simulation, depicted in Figure 3.12, takes two inputs, an individual's life history, generated by the natural history section, and a screening strategy. For this implementation, a screening strategy is defined by a screening test type, a screening frequency, and a starting and ending age for screening. The compliance rate for screening is incorporated into the model as a parameter that depends on the screening test employed. At the end of the screening component, a modified life history for the individual is generated by having the screening strategy to disrupt the natural progression of CRC within the individual. The screening timeline, for simplification purposes, is discretized in to one year increments, assuming that a screening test captures every event that happens within the year.

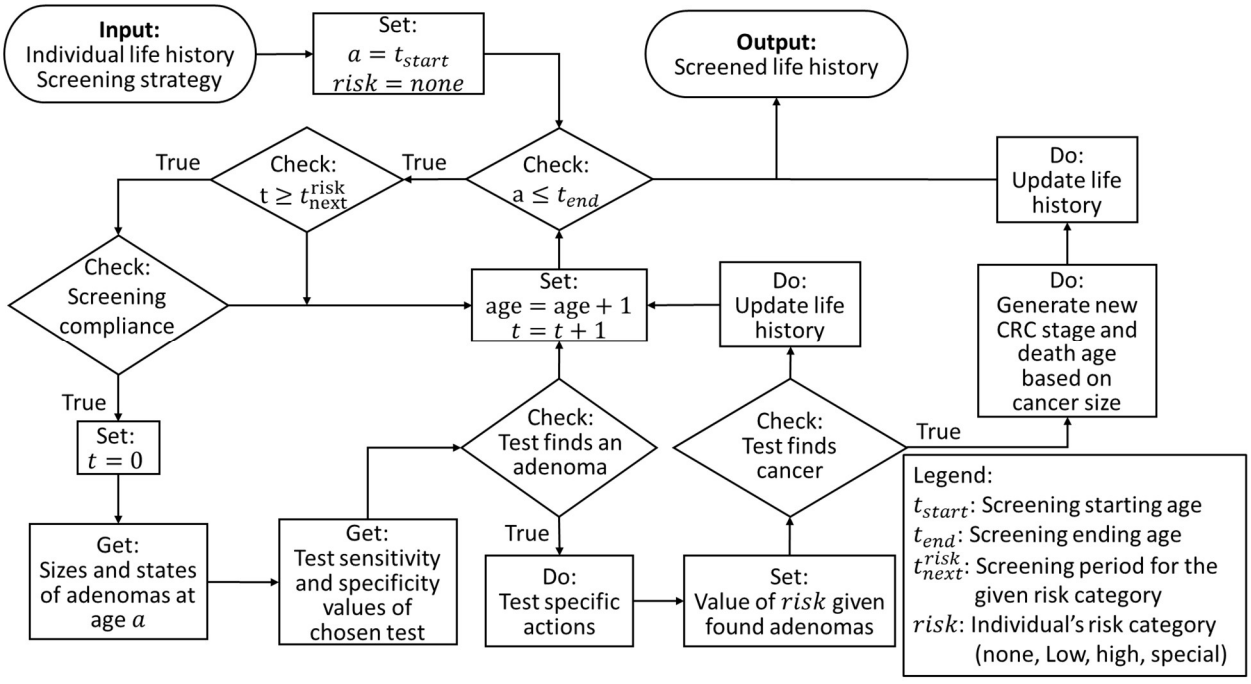

Figure 3.12 Flow chart of the screening and surveillance portion of the microsimulation.

There are two types of screening tests incorporated in the microsimulation: endoscopic and fecal tests. Each test has corresponding sensitivities and specificities associated with it[163], outlined

in Section 3.1.4. For endoscopic tests, colonoscopy and sigmoidoscopy are included. The only difference between the two comes from the reach of the test, where the colonoscopy can screen the entirety of the colon, but sigmoidoscopy can only reach to the end of sigmoid colon, or about 30 % of the entire colon[17]. For fecal tests, fecal immunochemical test (FIT) is included, where the test has a reach of the entire colon.

The screening component preforms two major actions depending on if the test is an endoscopic test or a fecal test. For the fecal test, the action starts with determining the age at which the individual has the test performed. From there, the number of adenomas the individual has at that age is evaluated. If there are no adenomas, a check for a false positive test result is performed via a random Uniform(0,1) value and comparing this value with the defined specificity of the test. For the visual tests, colonoscopy and sigmoidoscopy, the specificity is assumed to be 1.0, meaning there are no false positives. However, for the FIT the specificity is defined to be 0.93[19]. If the value is greater the specificity, a false positive test occurs. If at least one adenoma has developed within that person, a random Uniform(0,1) value is generated and compared to the sensitivity associated with that adenoma's current state, i.e., the size of the adenoma or if it has reached a cancerous state or not, providing a positive result if at least one adenoma, or cancer, is found. When a positive result is returned, a diagnostic colonoscopy is scheduled for that individual the same year, with the assumption that the individual is 100 % complaint with the diagnostic colonoscopy.

The endoscopic tests are the actions that actually modify an individual's life history. The initial evaluations are similar to that of the fecal test, determining the age of the individual and the presence of adenomas given the age. However, because we assume a perfect specificity there is no chance of a false positive. Each adenoma present at the time of screening is evaluated for detection based on the adenoma's state and location in the colon. If an adenoma is found prior to transitioning

to a cancerous state, a polypectomy is preformed, removing the adenoma from the individual. The adenoma removal also removes all remaining life events, such as transitioning to a cancerous state, associated with that adenoma from the individual's life history. Finally, if the endoscopic test being performed was a sigmoidoscopy and at least one adenoma was found, a diagnostic colonoscopy is scheduled with similar assumptions as in fecal test.

The screening component preforms these tests as defined by the input screening strategy. These strategies are limited to only allowing one screening test type to be chosen at a time, however this does not restrict the use of a diagnostic colonoscopy after positive test results. If a diagnostic colonoscopy or polypectomy is performed for any of the tested screening strategies, the individual is placed within a given risk category as defined in Table 3.5. The individual is then placed on surveillance screening with the time in between tests dependent on the risk category they were placed in. If the individual is diagnosed with cancer at any point within the screening process, they become ineligible for further screening. Once the individual reaches the end of screening age, an updated screened life history is output allowing a comparison to the natural life history of the individual to assess the effect of screening.

Table 3.5 Surveillance screening risk categories, definition, recommended range, and the used screening period [164,165]

Risk category	Category definition		Recommended range (years)	Screening period (years)
Low	1 – 2 adenomas (< 1cm)		5 – 10	7
High	3 -10 adenomas (< 1cm) OR 1 adenoma ≥ 1 cm	OR	3	3
Special	> 10 adenomas		1 - 3	2

3.1.4 Description and sources of model parameters

Four metrics, three associated with monetary costs and one based on QALY, are used to assess the effectiveness of a screening strategy. For each year of life, an individual has monetary cost and QALY values, which are discounted at an annual rate of 3%. The sum of the discounted values over the entire simulated cohort is utilized to assess the effectiveness of a given screening strategy. The monetary costs include screening and CRC treatments. The screening costs depend on the test type [166]. The treatment costs include the initial diagnosis, continuing treatment, and death associated with CRC [166]. The annual QALY uses a weight associated with the health state that the individual is in for a given year [167]. Additionally, there is a disutility associated with the endoscopic screening test modalities: a single decrement of 0.0055, or approximately two days, to that year's weight. A summary of the states and the associated weights and costs is in Table 3.6, alongside the costs for different screening tests.

Table 3.6 Annual QALY weights associated with CRC and related treatment costs, and screening costs

Health state	**Quality adjustment** [167]	**Cost for the year ($)** [166]
Healthy	1.0	0
Adenoma	1.0	0
Undiagnosed CRC	1.0	0
Diagnosed CRC stage I		
Initial detection	0.88	33,332
Continuing treatment	0.95	2,721
CRC death (year before)	0.3	61,312
Non-CRC death	0.95	15,106
Diagnosed CRC stage II		
Initial detection	0.82	46,330
Continuing treatment	0.95	2,537
CRC death (year before)	0.3	61,138
Non-CRC death	0.95	13,213
Diagnosed CRC stage III		
Initial detection	0.76	56,678
Continuing treatment	0.76	3,626
CRC death (year before)	0.3	64,421
Non-CRC death	0.76	17,480
Diagnosed CRC stage IV		
Initial detection	0.3	74,278
Continuing treatment	0.3	11,239
CRC death (year before)	0.3	86,458
Non-CRC death	0.3	46,934
Screening test	**Screening disutility**	**Cost per test ($)** [166]
FIT	None	30
Colonoscopy	0.0055	870
w/ polypectomy	0.0055	900
Sigmoidoscopy	0.0055	270
w/ polypectomy	0.0055	300

The detection rate of CRC and its precursors depends on the test used and the state of the disease. For a FIT, the detection rates for an adenoma < 10 mm, an adenoma ≥ 10 mm, and CRC are 6 %, 35 %, and 88 %, respectively [19]. For a colonoscopy, the CRC detection rate of CRC is taken to be 100 %, with the adenoma detection rate as defined in Eqn. (1) [52], where DR is the detection rate and s is the size of the adenoma in mm. A sigmoidoscopy uses the same detection rates as a colonoscopy. However, if the location of the adenoma or cancer is past the Sigmoid colon, the detection rate is 0 % regardless of stage.

$$DR = 1 - (0.34 - 0.035s + 0.0009s^2) \quad \textbf{3.15}$$

3.2 DFO Algorithm Determination

It has been shown in previous work[75,168] that the performance of a DFO solver can greatly depend on the problem needing to be solved. This section aims to determine the DFO solver that is best suited to solve the CRCSP, a test problem that shares similar characteristics to those of the CRC screening problem was constructed. The performance of ten DFO solvers that can handle combinatorial optimization problems, outlined in Section 2.2.1, were evaluated using this test problem and confirmed by using the CRCSP.

3.2.1 Evaluation of selected DFO Solvers on a Test Problem

3.2.1.1 Definition of the Test Problem

The DFO algorithms were used to solve several instances of a test problem that share the characteristics of the CRC screening problem. The number of decision variables was varied to explore the impact of problem dimensions on the performance of the solvers. The decision variables within the test problem are x_i, for $i = \{0,1, \dots, N\}$ where N depends on the test instance. The test problem instances, all nonlinear combinatorial programs, are given in Eqns. 3.16-3.22

(centered test problem), and in Eqns. 3.16-3.21 and 3.23 (off-centered test problem). They utilize a modified three dimensional, $N = 3$, (Eq. 3.22) and N-dimensional, $N = \{3,6,9\}$, (Eq. 3.23) Ackley function [169] as the objective (Eq. 3.16). The Ackley function is chosen due to its numerous local minima, a feature that is believed to be present within the CRC screening problem. Variable y_i, for $i = \{1,2, \dots, N\}$, is introduced to achieve symmetric solutions where the centered test problem has $N = 3$ and the off-centered test problem has $N \in \{3,6,9\}$. This variable (y_i) is calculated by rounding decision variable x_i to the nearest multiple of five (Eq. 3.17). Equation 3.18 defines the variable (a) that changes the objective function's maximum and minimum values. Similar to the screening problem, this variable can take on three values as restricted by the categorical variable x_0 (Eq. 3.21). Because the CRC screening problem is a combinatorial problem, the remaining decision variables x_i, for $i = \{1,2, \dots, N\}$, are restricted to integer domains, as shown in Eq. 3.22.

For visualization purposes, a two dimensional (2-D) continuous version of the objective function, $F(\boldsymbol{x})$ is rendered in Figure 3.13 (A) – (C) for the centered test problem with bounds presented in Eq. 3.22. The numerous local minima and the global minimum located at the center of the feasible region are the inherent features of the Ackley function. The plots (A), (B), and (C) of Figure 3.13 show the effect that the decision variable x_0 has on the shape of the objective function and the global minimum. As the value of x_0 increases from -1 to 1, the range of $F(\boldsymbol{x})$ also increases, with the greatest range and the smallest global minimum value corresponding to $x_0 = 1$, shown in plot C (Figure 3.13). The effects that variable y_i has on the symmetry of the problem is illustrated by the numerous flat regions within each plot of Figure 3.13. Finally, the bounds defined in Eq. 3.22 affect the location of the minima, with these bounds placing the minimum in the center of the feasible region.

$$\min F(\boldsymbol{x}) = -2a\exp\left(-b\sqrt{\frac{1}{3}\sum_{i=1}^{3}(y_i)^2}\right) - exp\left(\frac{1}{3}\sum_{i=1}^{3}\cos(cy_i)\right) + a + \exp(1) \quad \textbf{3.16}$$

$$s.t.\ \ y_i = 5\left\lfloor\frac{x_i}{5}\right\rfloor \qquad i = \{1,2,\dots,N\} \quad \textbf{3.17}$$

$$a = 20 + 10x_0 \quad \textbf{3.18}$$

$$b = 0.4 \quad \textbf{3.19}$$

$$c = \pi \quad \textbf{3.20}$$

$$x_0 \in \{-1,0,1\} \quad \textbf{3.21}$$

$$x_1, x_2, x_3 \in \{-32,-31,\dots,32\} \quad \textbf{3.22}$$

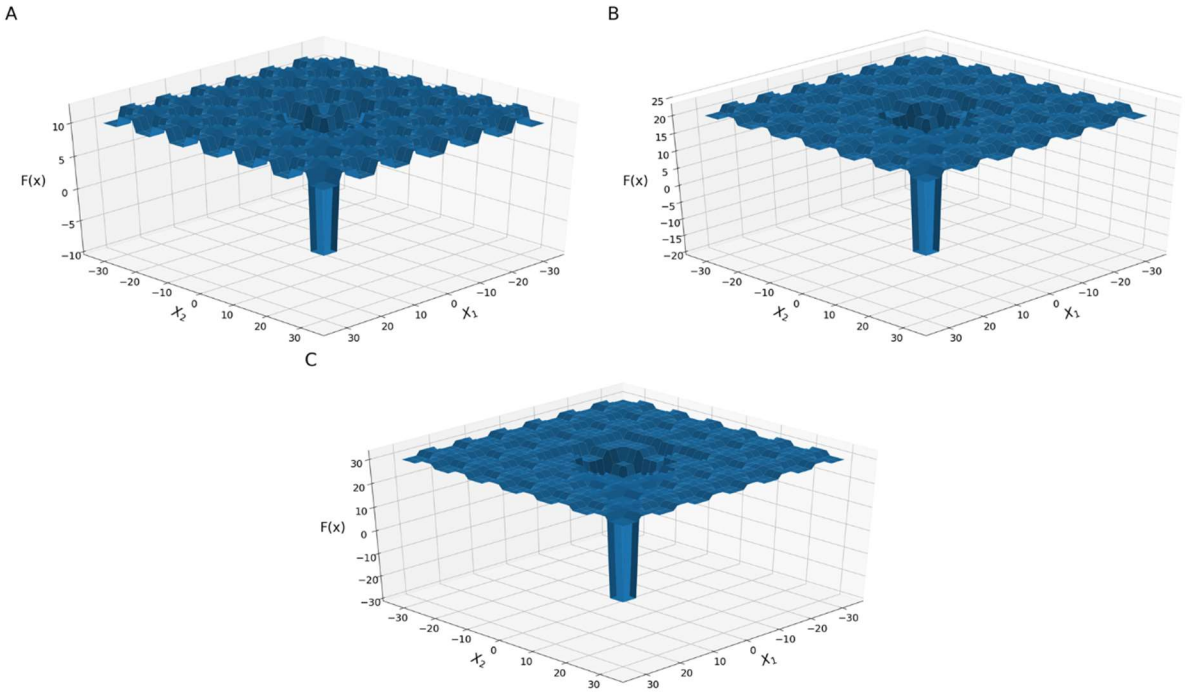

Figure 3.13 Continuous 2-D version of the centered test problem A) surface plot with $\boldsymbol{x_0} = -\mathbf{1}$, B) surface plot with $\boldsymbol{x_0} = \mathbf{0}$, C) surface plot with $\boldsymbol{x_0} = \mathbf{1}$.

Three instances of the off-centered test problem are constructed using the bounds given in Eq. 3.23 (replacing Eq. 3.22 of the centered problem). The global optima of these test instance are

near the boundary of the feasible region in contrast to the first instance where the global optimum is in the center of the feasible region. A rendering of 2-D continuous form of the off-centered test problem can be found in Figure 3.14 (A) – (C) for bounds consistent with Eq. 3.23. The only difference between Figure 3.13 and Figure 3.14 are the shift the bounds have on the location of the problem's global minimum.

$$x_1, x_4, x_7 \in \{-8, -7, \ldots, 56\}$$
$$x_2, x_5, x_8 \in \{-56, -55, \ldots, 8\}$$
$$x_3, x_6, x_9 \in \{-32, -31, \ldots, 32\} \quad \textbf{3.23}$$

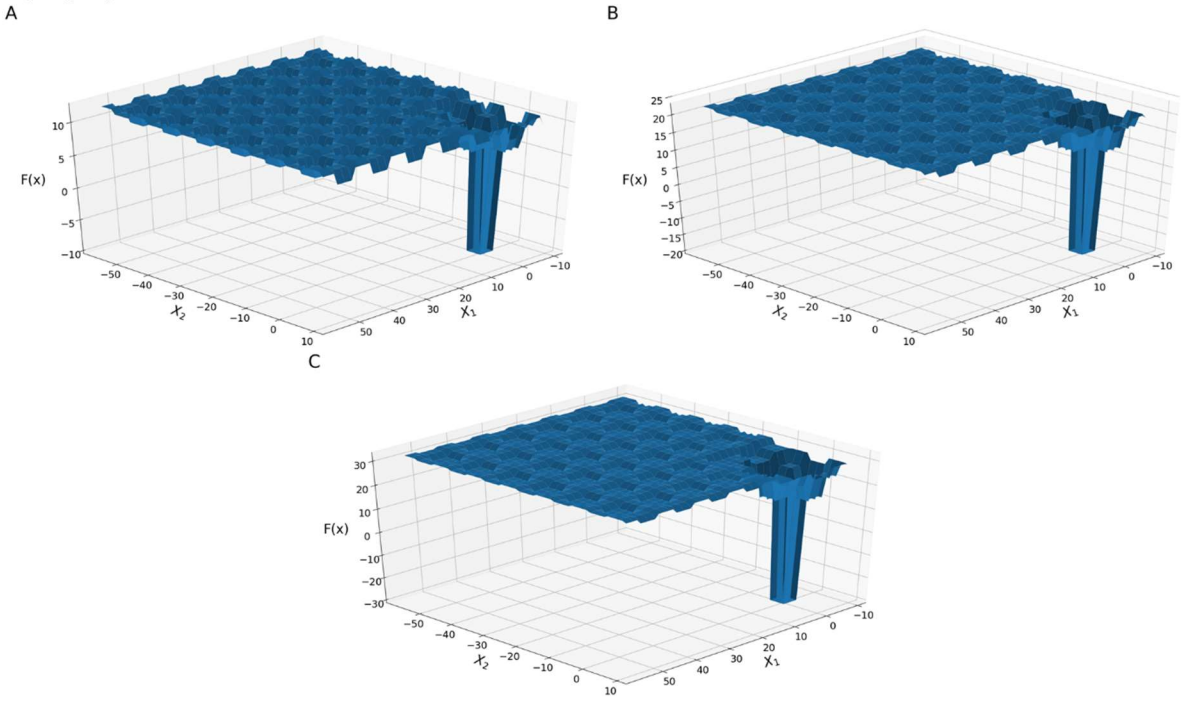

Figure 3.14 Continuous 2-D version of the off-centered test problem A) surface plot with $\boldsymbol{x_0 = -1}$, B) surface plot with $\boldsymbol{x_0 = 0}$, C) surface plot with $\boldsymbol{x_0 = 1}$.

3.2.1.2 Experimental Setup for the DFO Solvers Applied to the Modified Ackley Function Test Problem

The termination criterion for all DFO solvers was set to 1,000 model evaluations. The evaluation limit was chosen to provide a low computational budget but a large enough one to allow for the more efficient solvers to identify the optimum. For the domain given in Eq. 3.22, all solvers evaluated the randomly generated point $x_0 = 0, x_1 = 0, x_2 = 25, x_3 = 16$ as their initial point,

except for TOMLAB/glcSolve and TOMLAB/glcDirect where the initial point cannot be chosen and is set to $x_0 = 0, x_1 = 0, x_2 = 0, x_3 = 0$. For the domain given in Eq. 3.23, every solver evaluated $x_0 = 0, x_{1,j} = 24, x_{2,j} = -24, x_{3,j} = 0$ as the initial point, to provide each solver the same starting point as TOMLAB/glcSolve and TOMLAB/glcDirect. For solvers that require a set of initial points (instead of a single location), i.e., DEAP-GA, GPyOpt, MIDACO, CMA-ES, and RBFOpt, the same starting point was placed inside that initial set. If a DFO solver terminated before reaching 1,000 model evaluations, e.g., the local solvers may converge to a solution, a multi-start approach was used until 1,000 model evaluations were reached. The different starting points were generated randomly from a discrete uniform distribution covering the range of each decision variable.

Default algorithmic parameters and settings were used for each solver, except for DEAP-GA and SimAnneal solvers, which require a set of user-defined parameters. The default parameters are used to evaluate the out-of-the-box capabilities of the solvers. Table 3.7 summarizes the versions of the solvers used within the study. For DEAP-GA solver, the population per generation was set to 50, the next generation was selected via a roulette process, the mating method used a 2-point crossover with a probability of occurrence of 50%, and the mutation rate was set to 25%. The SimAnneal solver used the default settings for starting temperature and ending temperature, 25,000 and 2.5, respectively. The movements between states were defined such that each value for x_0 can be chosen with equal probability for the new state, and the step size for the remaining variables was set via random sampling from a Uniform(-16,16) distribution, corresponding to plus or minus a quarter of the entire range of the variable's bounds. If the step size were to exceed a decision variable's bounds at any state, the step size was adjusted to the bound that is exceeded.

Table 3.7. Solver versions used within the study

Solver	Version	Solver	Version
SimAnneal	0.4.2	CMA-ES	3.33.integer
DFL	NA	TOMLAB/glcSolve	8.5
DEAP – GA	1.2.2	RBFOpt	4.0.2
MIDACO	6.0	NOMAD	3.8.1
TOMLAB/glcDirect	8.5	GPyOpt	1.2.5

For solvers that allowed decision variables to explicitly be defined as categorical variables, i.e., NOMAD and GPyOpt, x_0 was defined as a categorical variable. For the rest, three binary variables, $x_{0,1}, x_{0,2}$ and, $x_{0,3}$, were defined where a value of $x_{0,1} = 1$ represents $x_0 = -1$, $x_{0,2} = 1$ represents $x_0 = 0$, and $x_{0,3} = 1$ represents $x_0 = 1$. A new constraint, given in Eq. 3.24, was added to the optimization problem instances. For the solvers that can only handle unconstrained optimization problems, a null evaluation was returned as the objective function value to the solver when this constraint (Eq. 9) was violated. The RBFOpt solver could not handle a null value, instead, an output of 10^{12}, a large enough value compared to the objective function values to represent infinity, was returned. The implementation of MIDACO solver had a limit on the number of decision variables it can handle, therefore, x_0 was treated as a categorical variable. The test problems were solved twice using TOMLAB solvers, glcSolve and glcDirect, once x_0 defined as a categorical variable and once using the binary variables ($x_{0,1}, x_{0,2}$ and, $x_{0,3}$) with the additional constraint, Eq. 3.24. The DFL solver was only used to solve the first instance of the test problem twice, first using categorical then using the binary decision variable definitions. Based on the results, it was concluded that the decision variable x_0 should be defined as a categorical variable for the second instance of the test problem for the DFL solver. The test problem instances were implemented in Python 3.5.

$$x_{b,1} + x_{b,2} + x_{b,3} = 1 \quad \textbf{3.24}$$

Both NOMAD and DFL solvers were run using compiled executables, accessing Python 3.5 for the test problem. The CMA-ES and TOMLAB solvers were run using MATLAB R2015b, accessing Python 3.5 for the problems. The remaining solvers were implemented in Python 3.5, all on an Intel Xeon E5-2650 2.30 GHz processor.

Two initial criteria were used to evaluate the performance of the solvers. Three additional criteria were defined to compare the local search solver performances. The top-performing solvers were evaluated using two additional criteria. The first criterion for all solvers was the absolute percent difference from the optimum (Eqn. 3.25) and how it changes as a function of the total number of objective function evaluations.

$$\%\, diff = \left| \frac{y_i^* - y^{opt}}{y^{opt}} \right| \quad \textbf{3.25}$$

In Eq. (10), y_i^* is the current best-identified solution by the solver for function evaluation i, and y^{opt} is the global minimum of the function. The second criterion was the solution time, which includes the time the solver took to generate and perform the 1,000 function evaluations.

For the local search solvers, the first additional performance criterion was the percentage improvement given a starting point, given by Eqn. 3.26,

$$\%\, improvement = \frac{y^* - f(\boldsymbol{x}_0)}{y^{opt} - f(\boldsymbol{x}_0)} \quad \textbf{3.26}$$

where $f(\boldsymbol{x}_0)$ is the function value at the starting point $\boldsymbol{x}_0$, y^* is the best-identified solution by the solver using the starting point $\boldsymbol{x}_0$, and y^{opt} is the global minimum of the function. The second and third additional criteria for local solvers are the number of times $\%\, improvement$ was equal to one and the number of times $\%\, improvement$ was equal to zero for the various starting points tested. Finally, the criteria to evaluate the top performers are the percentage of

instances in which the solver identified the global minimum and the number of function evaluations needed to identify the global minimum.

3.2.1.3 Results of the Modified Ackley Function Tests

3.2.1.3.1 Comparison of the Solution Times

The times taken by solvers to generate and evaluate the 1,000 functional calls are summarized in Table 3.8. The results in Table 3.8 reveal that the solution times were comparable to each other for all solvers, aside from three outliers RBFOpt, NOMAD, and GPyOpt, with GPyOpt being the most extreme. The outlier solvers construct an approximate model of the objective function and solve an optimization problem using it to generate the next point to evaluate, which in turn takes more time than direct search algorithms. The differences in solution times of RBFOpt, NOMAD, and GPyOpt are a direct outcome of the efficiency of each solver in generating the approximate model and in solving the resulting optimization problem. The approximate model building times are a significant portion of these solvers' run time and make them less efficient than the direct search solvers when the evaluation of the black-box model is not computationally expensive. However, when the objective function becomes more computationally expensive to evaluate, the time disparity between model-based and direct search algorithms has less of an impact on the total solution time.

Table 3.8. Evaluated DFO solution times (wall time) for the centered test problem (total of 1000 function calls).

Solver	Wall Time (h:mm:ss)	Solver	Wall Time (h:mm:ss)
SimAnneal	0:00:04	CMA-ES	0:00:05
DFL	0:00:04	TOMLAB/glcSolve	0:00:06
DEAP – GA	0:00:05	RBFOpt	0:03:22
MIDACO	0:00:05	NOMAD	0:07:59
TOMLAB/glcDirect	0:00:05	GPyOpt	4:37:05

3.2.1.3.2 Comparison and Analysis of the Local Search Solvers

The solvers DFL and NOMAD were the local ones. Figure 3.15 shows the absolute percent difference from the optimum as a function of total black-box evaluations for the centered test problem. The NOMAD solver utilizes random numbers within the search step of the solver; therefore, we performed a series of 50 runs, each reaching 1,000 function evaluations, to assess its performance. In Figure 3.15, the line denoted "Best NOMAD" shows the run that converged on to the optimal solution the fastest, "Worst NOMAD" shows the run that, after the 1,000 evaluations, had the worst gap, and "Median NOMAD" shows the median case of the 50 runs. The median case is determined by evaluating the median value of the absolute percent difference from the optimum. Suppose multiple runs end with the same value of this criterion. In that case, we evaluate the area under the absolute percent difference curve as a function of evaluations and take the run with the median value. For the DFL solver, it was found that the performance was superior when model implementation used a categorical variable, both in terms of locating the optimum as well as general search direction rather than implementing the model using a set of binary variables. The improvement is due to DFL perturbing a single variable at a time, therefore violating the constraint

in Eq. 3.24 when attempting to vary the binary variables. This causes DFL never to change the set of binary category variables and instead only to stay on the starting value.

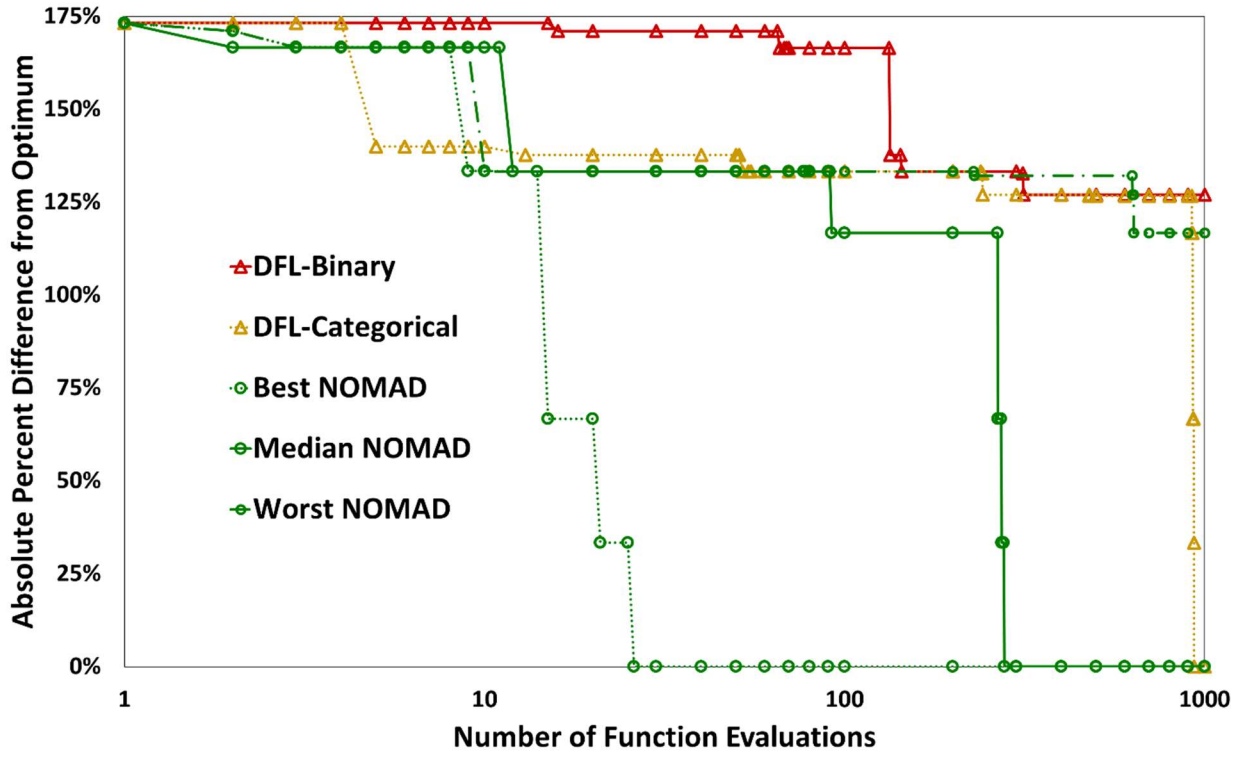

Figure 3.15 Change in absolute percent difference from the optimum with the number of function evaluations for the local search solvers (DFL and NOMAD) on the centered test problem with the domain of: $\boldsymbol{x_1, x_2, x_3} \in \{-\mathbf{32}, -\mathbf{31}, \dots, \mathbf{32}\}$.

The percent improvements obtained from 30 different starting points by each solver are compiled in Table 3.9, where the average improvement by each solver, number of times the solver found the optimum, and the number of times the solver failed to yield any improvements are presented. The metrics for NOMAD report the median per each starting point, with the 2.5 and 97.5 percentiles included in parenthesis. The results in Table 3.9 support the previous conclusions that DFL is more effective using a categorical variable than a set of binary variables for solving the test problem, yielding about a 15% increase on the average improvement and halving the number of times DFL was unable to improve from the starting point. Overall, the performance of

NOMAD is better than that of DFL for all criteria in Table 3.9. Average improvement over 30 different starting points for NOMAD and DFL for the centered test problem. The large range of results observed for NOMAD (20.6% to 43.5% improvements at 2.5 and 97.5 percentiles) prompted the question of whether the starting point or the random seed had a more significant effect. We calculated the variance of the average improvement for each starting point based on the random seed and compared it with the variance of the average improvement for each seed based on the starting points. The variance for each starting point was 0.032, while the variance for each random seed was 0.004. These results suggest that the starting point has a more significant effect on the performance of NOMAD compared to the random seed.

Table 3.9. Average improvement over 30 different starting points for NOMAD and DFL for the centered test problem

	NOMAD	DFL Categorical	DFL Binary
Average improvement	21.5 % (20.6%, 43.5%)	17.8 %	3.1 %
Number of times optimum located	5 (1.45, 9)	1	0
Number of times with no improvements	2 (1, 4)	5	10

The absolute percent difference from the optimum per function evaluation for the off-centered test problem is presented in Figure 3.16, with the local solver specific evaluation criteria in Table 3.10. Based on the comparative analysis of the DFL's performance on both model implementations of the centered test problems, only the model implementation using the categorical variable is included for the off-centered test problem. As can be seen from Figure 3.16, DFL performs well, finding the optimum on the first starting point. However, from Table 3.10, we

note that DFL was only able to find the optimum once out of all 30 starting points. Table 3.10 reveals a notable increase in the average improvement for the off-centered problem when compared to the centered problem (Table 3.9). This improvement, however, may be a result of the starting points generated for the new bounds. On average, the objective function values of the new starting points were 10% greater than the ones obtained for the centered test problem giving the two solvers a larger opportunity to improve the objective function values. Overall, our findings show that DFL is likely to converge on to a local optimum. However, if the starting point for DFL is close to, or on a local optimum, DFL generally fails to improve that solution for functions that have numerous local minima with many symmetric solutions. Alternatively, NOMAD has the ability to significantly improve a solution that is close to a local minimum, but drastic improvement is not always guaranteed for these types of problems.

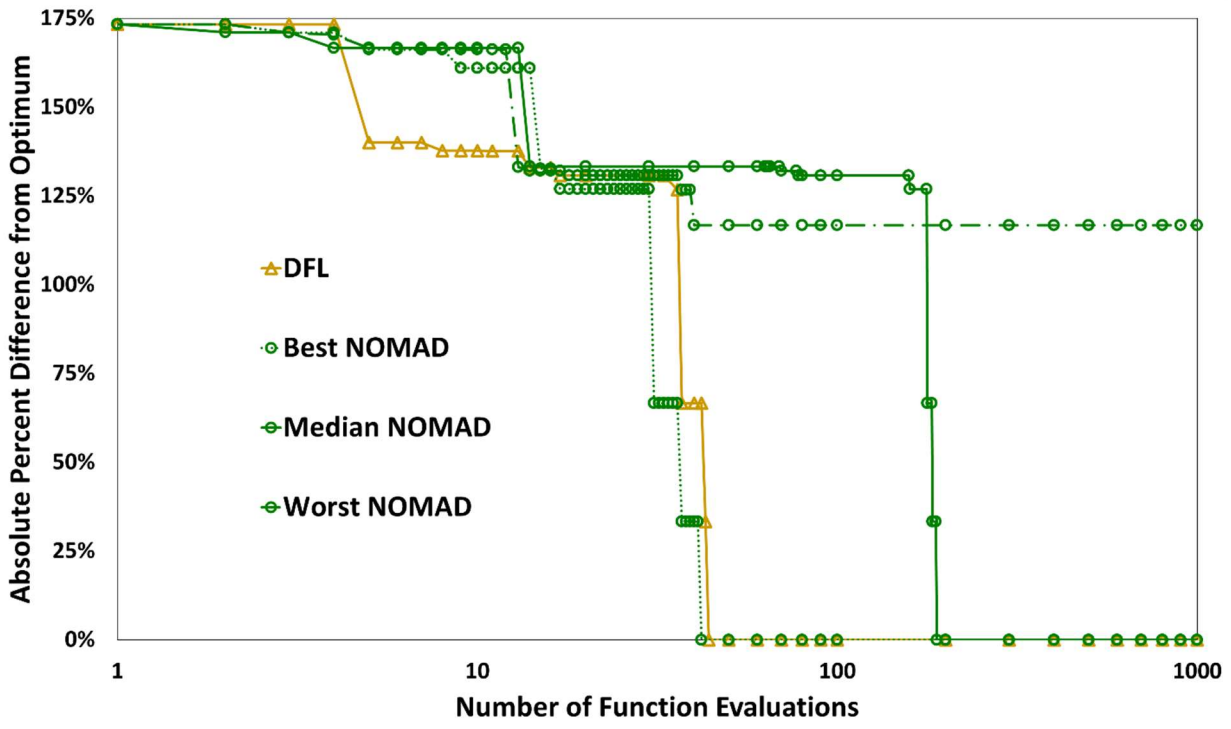

Figure 3.16 Change in absolute percent difference from the optimum with the number of function evaluations for the local search solvers (DLF and NOMAD) on the off-centered test problem with the domain of: $\boldsymbol{x_1} \in \{-\mathbf{8}, -\mathbf{7}, \dots, \mathbf{56}\}, \boldsymbol{x_2} \in \{-\mathbf{56}, -\mathbf{55}, \dots, \mathbf{8}\}, \boldsymbol{x_3} \in \{-\mathbf{32}, -\mathbf{31}, \dots, \mathbf{32}\}$.

Table 3.10 Average improvement over 30 different starting points for NOMAD and DFL for the off-centered modified Ackley function

	NOMAD	DFL
Average improvement	37.1 % (29.4%, 51.1%)	26.0 %
Number of times optimum located	5 (2, 10.8)	1
Number of times with no improvements	3 (1.2, 5)	4

3.2.1.3.3 Comparison and Analysis of the Global Search Solvers

Global search solvers have the tendency to rely heavily on random values to search the entire domain. As such, four out of six of the global search solvers tested fall into the stochastic search category and were, therefore, run for 50 times. TOMLAB/glcSolve and TOMLAB/glcDirect, however, were only run once due to their deterministic nature. Figure 3.17 plots the change in absolute percent difference from the optimum with the number of function evaluations for each solver. In Figure 3.17, the values observed for solvers SimAnneal, DEAP-GA, CMA-ES, and GPyOpt are the median values of the 50 runs of the respective solver.

Of the stochastic solvers, GPyOpt was the only solver that consistently found the optimal value within the allotted number of function evaluations. Interestingly and contrastingly to other studies [75,168], solver CMA-ES yielded the worst average performance, which we speculate may be due to the number of symmetric solutions for the test problem. The symmetric solutions create flat regions within the problem in a near checkerboard manner; see Figure 3.13Figure 3.14. The function values in these flat regions change abruptly, creating a jagged surface. The CMA-ES relies on updating its covariance matrix with the objective values from the set of all trial points to effectively explore the solution space. If most, or all, of the trial points share a similar objective value, no new information is available CMA-ES for update the covariance matrix. Hence, the

CMA-ES solver does not have additional information regarding the direction of the global minimum and generates random samples for the search, similar to the behavior of SimAnneal.

The comparison between TOMLAB/glcSolve and TOMLAB/glcDirect and the remaining solvers are of little value for the centered test problem because TOMALB/glcSolve and TOMLAB/glcDirect start their search at the center of the feasible region, which is very close to the global optimum. However, the two solvers and the different decision variable implementations of the DIRECT algorithm can be evaluated against each other. For the DIRECT algorithm, it is noted that, due to the curse of dimensionality, efficiency suffers as the number of variables increases [91]. This can be seen through the comparison of glcSolve-bin and glcSolve-cat, as the glcSolve-cat, where the model utilizes the categorical variable, identifies the global optimal much quicker than the glcSolve-bin where binary variables were employed. It should be noted that for the model with the categorical variable, glcSolve and glcDirect behaved identically, following the same solution path. A difference is noticed, however, when comparing the solution paths for the model with the binary variables. Figure 3.17 reveals that glcSolve locates the global minimum in nearly $1/10^{th}$ of the total function evaluations than it took for glcDirect to find the global minimum. glcDirect attempted to evaluate points that violated the equality constraint associated with the binary variables (Eqn. 3.24) during the initial function evaluations suggesting potential difficulties in the handling of equality constraints.

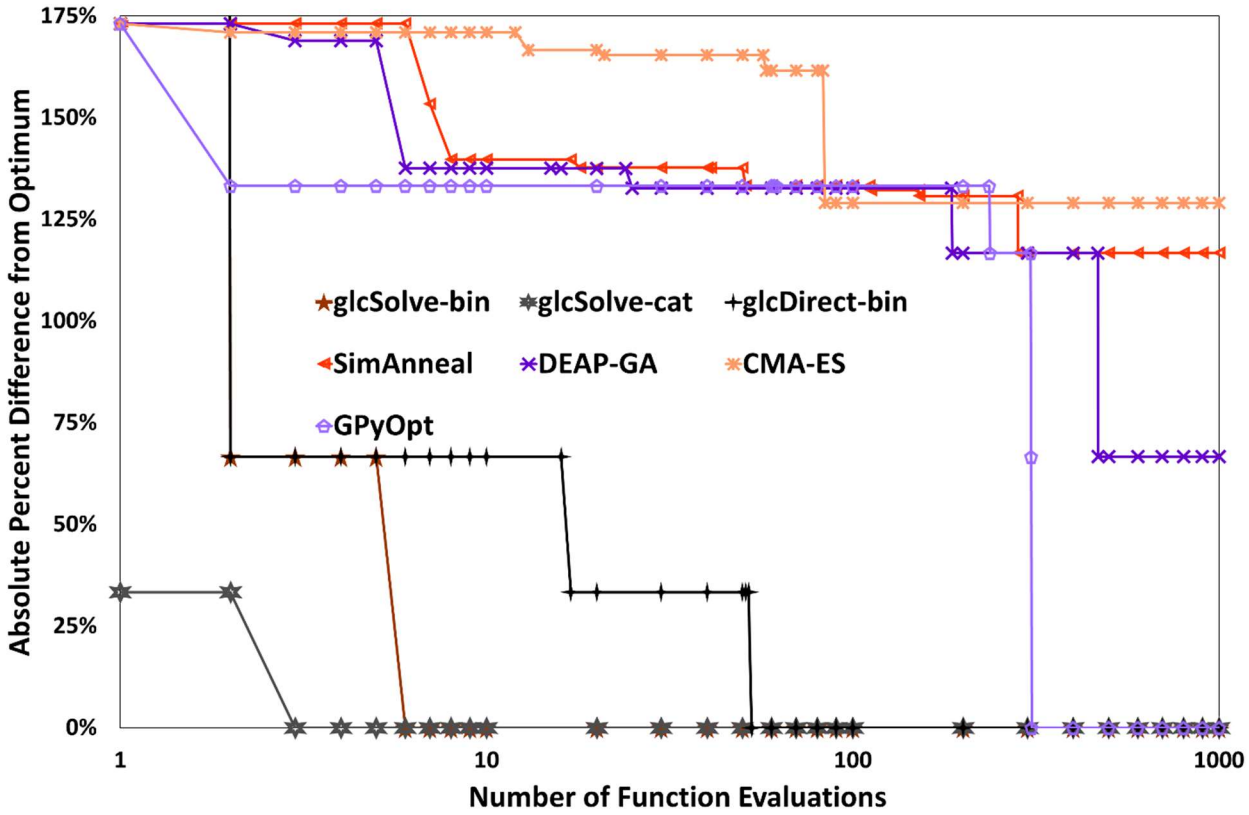

Figure 3.17 Change in absolute percent difference from the optimum with the number of function evaluations for the global search solvers (glcSolve-bin & cat, glcDirect-bin, SimAnneal, DEAP-GA, CMA-ES, and GPyOpt) on the centered test problem with the domain of: $\boldsymbol{x_1, x_2, x_3} \in [\boldsymbol{-32, -31, \ldots, 32}]$.

The DFO performances in terms of absolute percent difference for the off-centered test problem can be seen in Figure 3.18, with the absolute percent difference from the optimum for the stochastic search based solvers corresponding to the median value of the 50 runs. The two TOMLAB solvers outperform the remaining global search solvers. GPyOpt was the only other solver that found the solution at least half the time. For SimAnneal and DEAP-GA, there was very little change between the two versions of the test problem. SimAnneal was not able to improve the initial solution considerably on either version, and DEAP-GA converged to the same solution for both problems. On the other hand, CMA-ES located a solution with 33 % difference for the off-centered problem compared to 129 % for the centered one, a significant improvement. The difference could be due to the location of the global optimum, in the corner of the search space for

the off-centered problem (see Figure 3.14). The global optimum is located in a region with a steep decrease in objective value, and the region surrounding the optimum is equal in objective value (see Figure 3.13). Due to this, the likelihood of the solver overshooting the optimum and finding similar solutions is decreased by the reduction of symmetric solutions surrounding the optimum from shifting the bounds. Similar arguments could be made regarding GPyOpt performance; if the acquisition function for GPyOpt were to sample from the region around the global minimum, it would be sampling the points of the same value. These samples could suggest that instead of a sudden decrease in values, a flat region would be present instead. This could drive the acquisition function to sample the local minima rather than sampling what could seem like an entire region of poor-quality solutions.

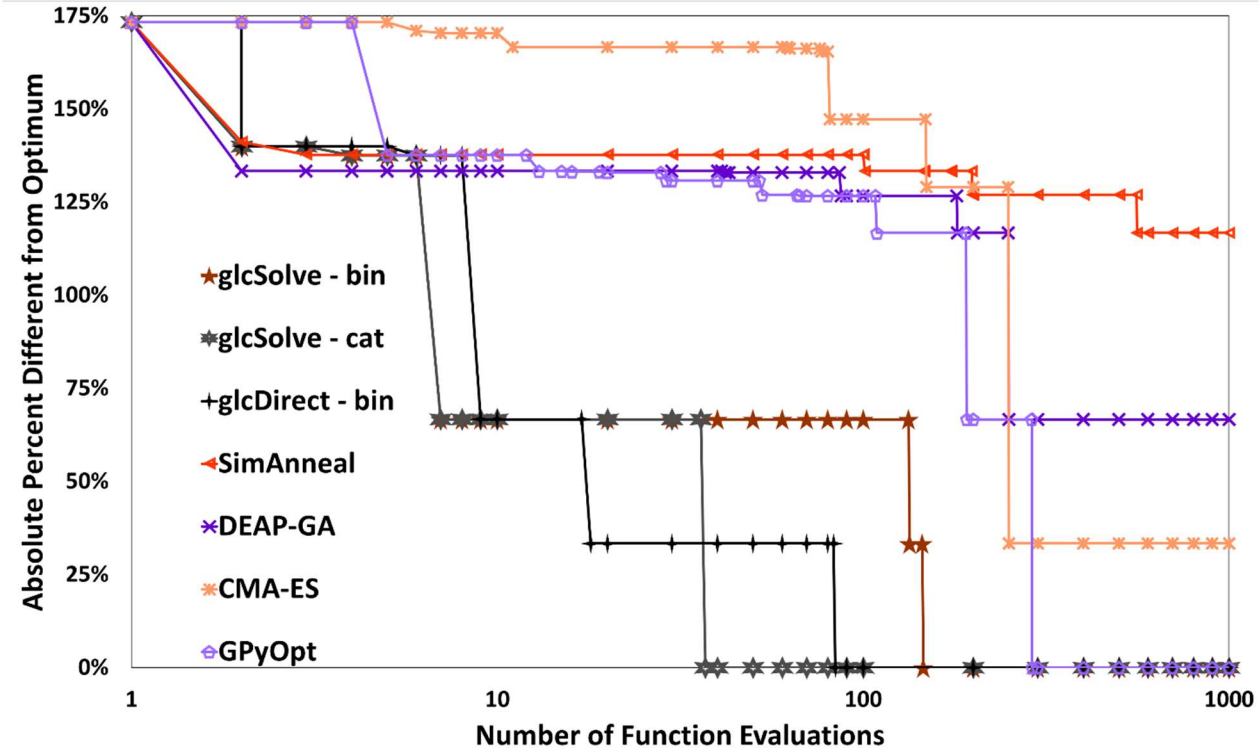

Figure 3.18 Change in absolute percent difference from the optimum with the number of function evaluations for the global search solvers (glcSolve-bin & cat, glcDirect-bin, SimAnneal, DEAP-GA, CMA-ES, and GPyOpt) on the off-centered test problem with the domain of: $\boldsymbol{x_1} \in \{-8, -7, \ldots, 56\}, \boldsymbol{x_2} \in \{-56, -55, \ldots, 8\}, \boldsymbol{x_3} \in \{-32, -31, \ldots, 32\}$.

3.2.1.3.4 Comparison and Analysis of the Hybrid Search Solvers

A plot presenting the absolute percent difference from the optimum versus the total number of black-box evaluations can be found in Figure 3.19 for the first test problem. For both RBFOpt and MIDACO, 50 runs were performed. The best run is marked with a dotted line, the worst is marked with a dashed line, and the median of the runs is marked with a solid line in Figure 3.19. The same plot for the second test problem is given in Figure 3.20.

Both figures reveal that, on average, MIDACO finds the optimum quicker than RBFOpt. Because RBFOpt constructs a surrogate model, it generally requires more function evaluations to build an accurate enough model to identify the region where the optimum lies. For both test problems, RBFOpt was able to provide a better solution in the worst case. However, those solutions still had a substantial gap from the true optimum. The categorical variable in the problems was represented using a set of binary variables for the model solved by RBFOpt as per its requirements. RBFOpt cannot solve constrained optimization problems. Hence, a very large objective function value was returned for infeasible functions calls (i.e., when Eq. 3.24 was violated) generated by RBFOpt, which potentially resulted in a poor fit surrogate model around the neighborhood of the points that violate the constraint. Shifting of the bounds for the problem had no substantial effect on either of the solvers' performance, and they were both able to reliably find the optimum solution. Across the 50 runs, MIDACO was able to find the global optimum 38 times for both test problems. For RBFOpt, there was a slight increase in performance for the off-centered problem compared to the centered one, from 39 runs to 43 runs in which the global optimum was found.

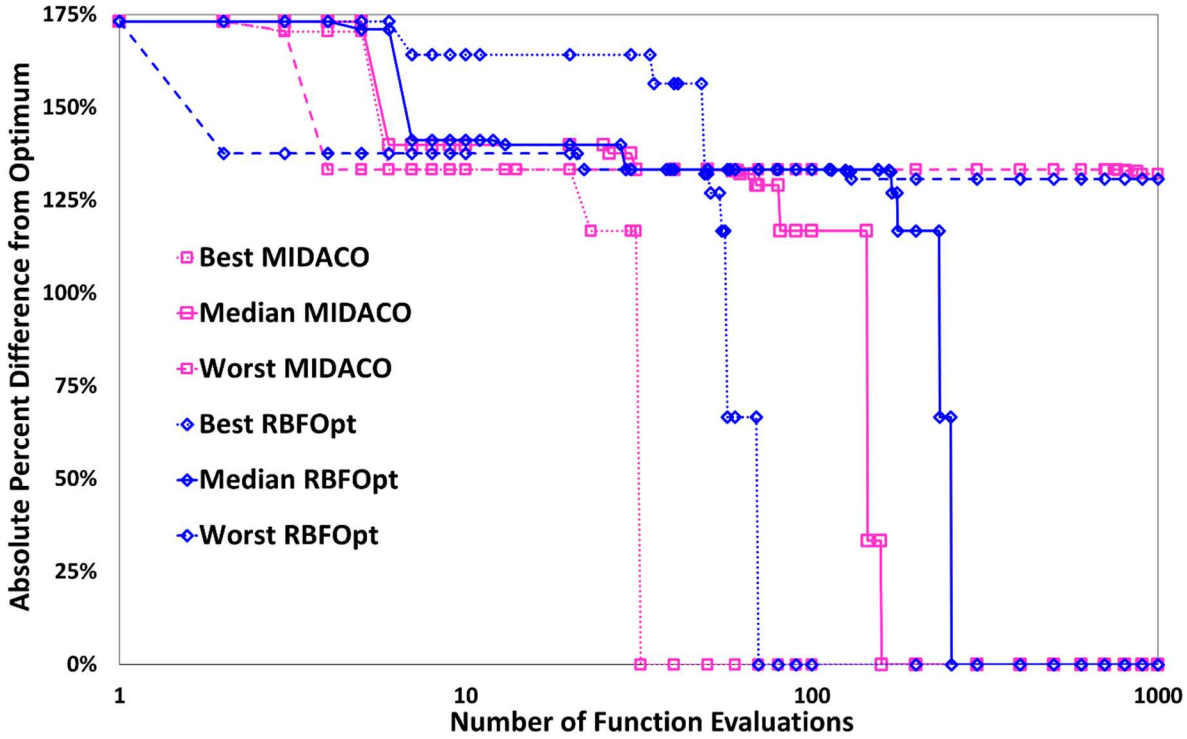

Figure 3.19 Change in absolute percent difference from the optimum with the number of function evaluations for the hybrid search solvers (MIDACO and RBFOpt) on the centered test problem with the domain of: $\boldsymbol{x_1, x_2, x_3 \in [-32, -31, \ldots, 32]}$.

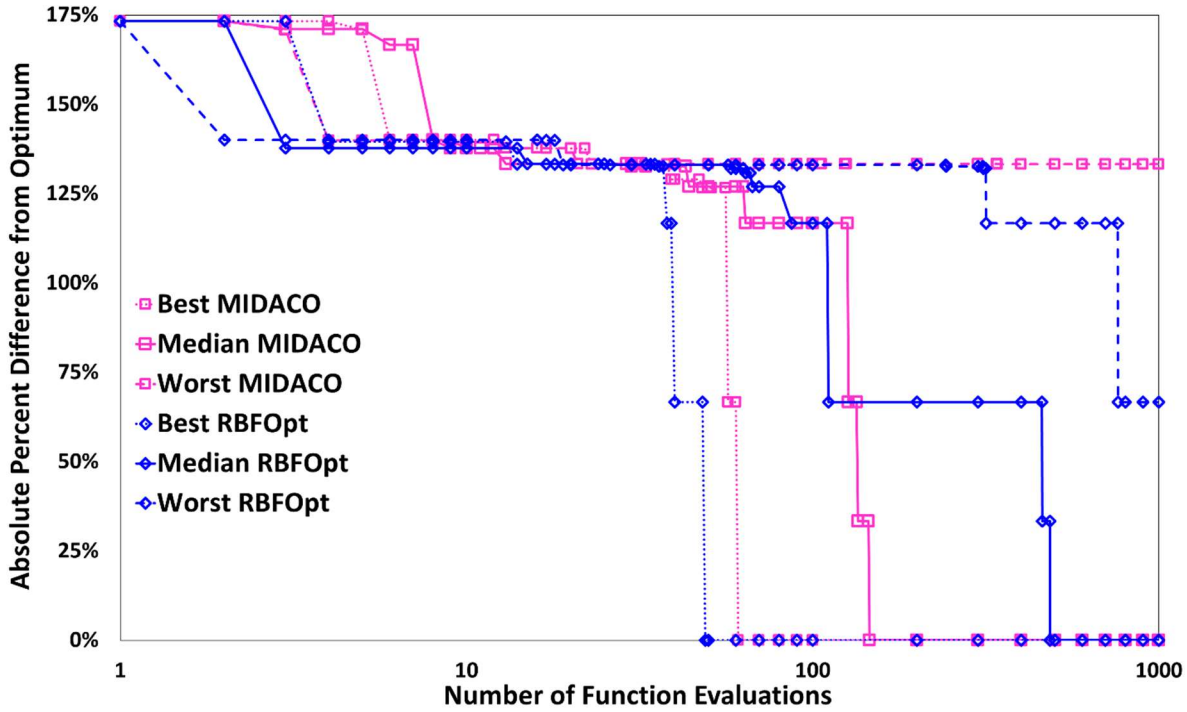

Figure 3.20 Change in absolute percent difference from the optimum with the number of function evaluations for the hybrid search solvers (MIDACO and RBFOpt) on the off-centered test problem with the domain of: $\boldsymbol{x_1 \in \{-8, -7, \ldots, 56\}, x_2 \in \{-56, -55, \ldots, 8\}, x_3 \in \{-32, -31, \ldots, 32\}}$.

3.2.1.3.5 Effect of Problem Dimension

Figures Figure 3.21, Figure 3.22, andFigure 3.23 present the absolute percent difference from the optimum versus the number of function evaluations for the four, six, and 10 variable variants of the off-centered test problem. The values used in the plots for the solvers that utilize random values (CMA-ES, DEAP-GA, GPyOpt, MIDACO, NOMAD, RBFOpt, and SimAnneal) use the median of 50 different runs to assess the solvers' average performance.

The solver DFL exhibited the best performance with the increase in the number of variables. As with the 4-variable off-centered problem, DFL was able to identify the optimum on the first starting point for the 7- and 10-variable versions of the off-centered problem. However, the same trend for DFL was observed in that the first starting point was the only starting point from which the solver converged to the optimal solution. Though none of the other solvers were able to identify the optimum on average as the variables increased, some solvers were able to identify the optimal solution a percentage of the time. For the seven-variable problem, RBFOpt, NOMAD, CMA-ES, and MIDACO were able to identify the optimum 30 %, 26 %, 12 %, and 6 % of the time, respectively. For the 10-variable problem, only NOMAD was able to identify the optimum more than once at 12 % of the time. The solvers, MIDACO and DEAP-GA, were the only other solvers that identified the optimum once out of the 50 runs.

The performance of each solver was expected to worsen with an increase in the number of variables while leaving the same computational budget in place. This trend was observed for each of the solvers. The decline in performance for both GPyOpt and RBFOpt can be attributed to the increase in solution space, creating a need for a higher number of function evaluations to construct a reliable approximate model. However, the computational budget did not allow these solvers to identify the optimal solution's neighborhood correctly. As mentioned in Section 3.2.1.3.3, the decline of performance for the glcSolve and glcDirect can be attributed to the DIRECT algorithms

noted decline in performance as the variables increase [91]. Solvers that use evolutionary methods, DEAP-GA, SimAnneal, CMA-ES, and MIDACO, also suffer as their ability to search the solution space effectively is severely hampered by the computational budget. Out of the four, DEAP-GA performed the best on average as the number of variables increased. Although, both MIDACO and CMA-ES performed better on the 7-variable problem in the identification rate of the optimum. Although solvers, DFL and NOMAD, exhibited diminishing performance with increasing problem dimensionality, their performance was the most robust. We hypothesize that because these solvers utilize local search that converges relatively quickly, the multi-start approach allowed for the convergence to multiple local minima within the allotted computational budget.

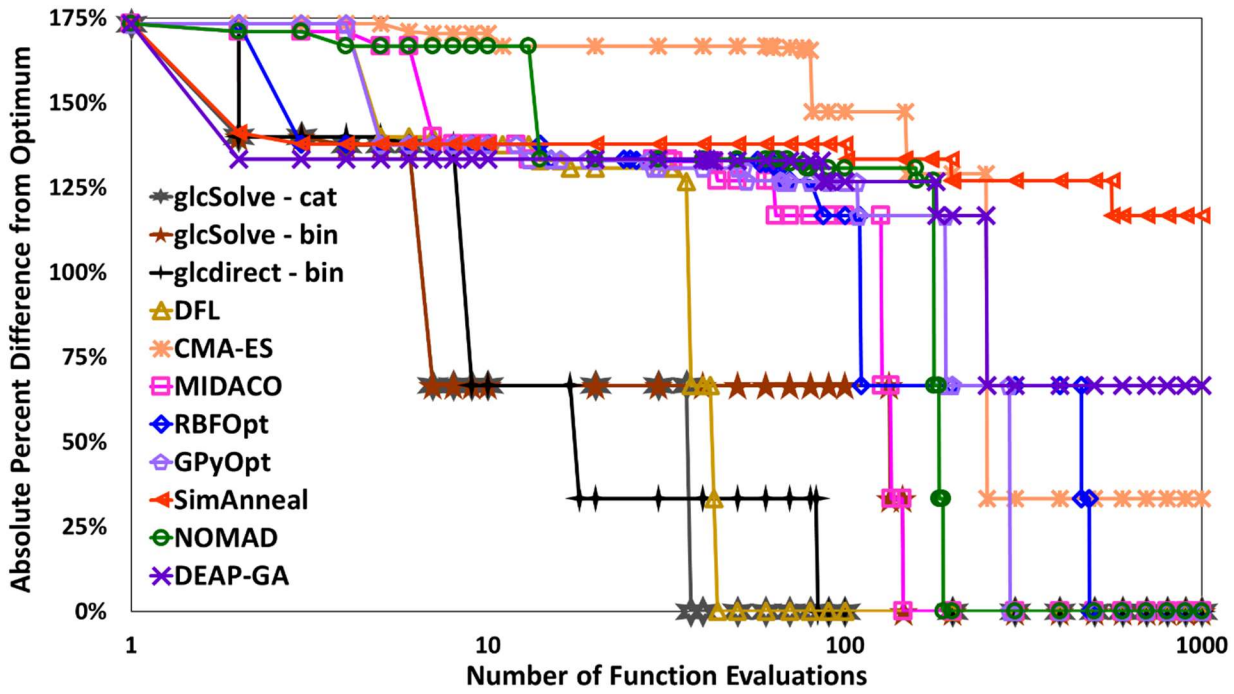

Figure 3.21 Change in absolute percent difference from the optimum with the number of function evaluations for all solvers on the 4-variable off-centered test problem.

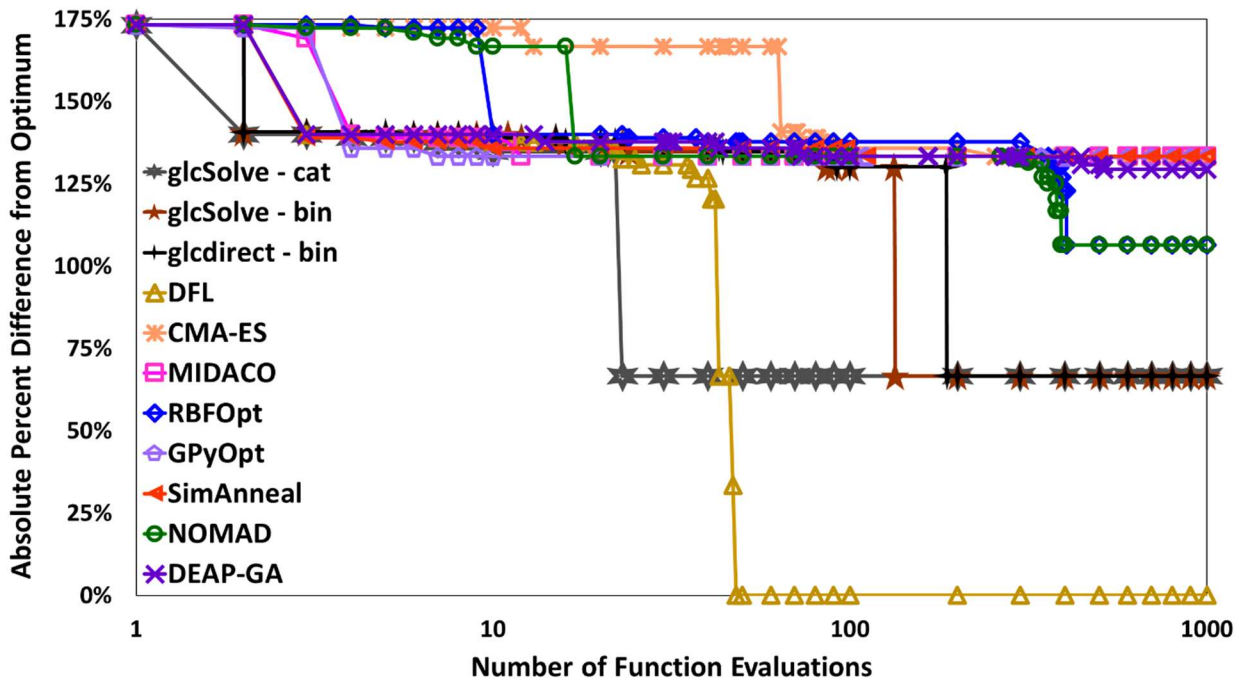

Figure 3.22 Change in absolute percent difference from the optimum with the number of function evaluations for all solvers on the 7-variable off-centered test problem.

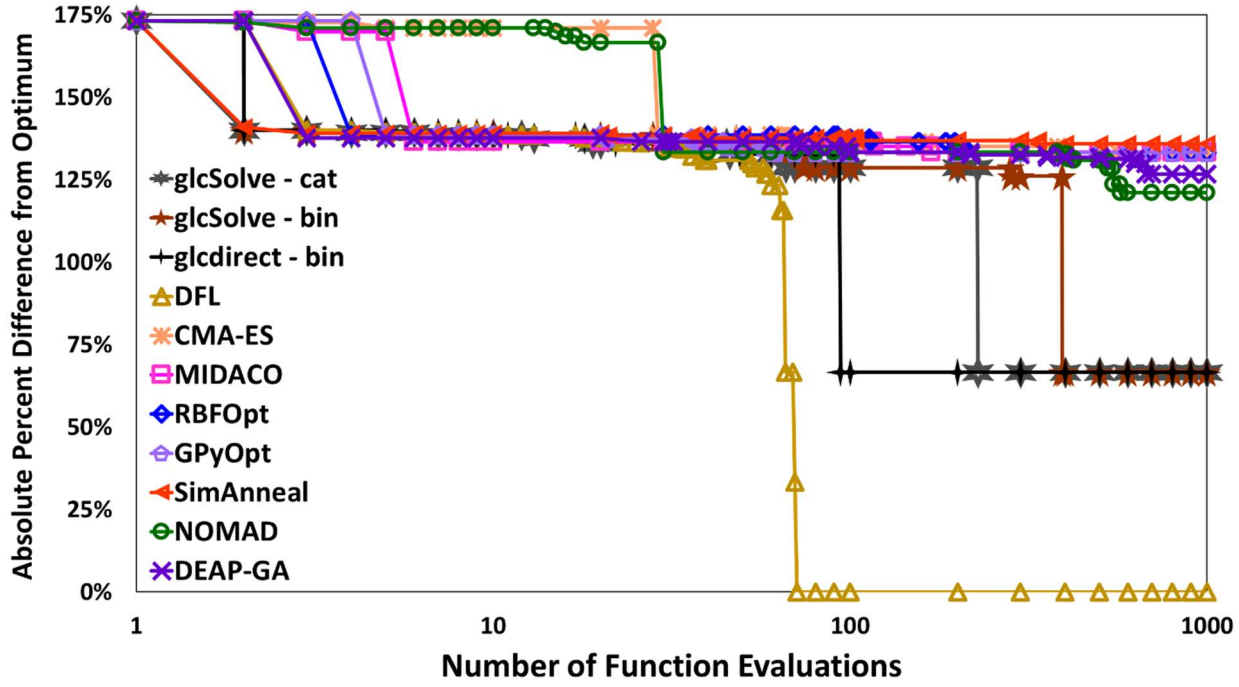

Figure 3.23 Change in absolute percent difference from the optimum with the number of function evaluations for all solvers on the 10-variable off-centered test problem.

3.2.1.3.6 Comparison of the Best Performing Solvers

Solvers were considered the best performers if they were able to identify the global optimum at least 50% of the time for the problems with four variables, most similar to our CRC screening test problem. These solvers are DFL, NOMAD, GPyOpt, TOMLAB/glcSolve, and RBFOpt. Given the performances of TOMLAB/glcSolve and TOMLAB/glcDirect when solving both test problems using the categorical and binary variables, TOMLAB/glcSolve was selected for this comparison. Table 3.11 outlines the number of evaluations in which the solver identified the global minimum, the percentage of runs where the solver located the optimum, and the number of invalid function evaluations. For solvers other than DFL and TOMLAB/glcSolve, the values in parenthesis are the best and worst cases, where the global minimum was still identified. DFL and TOMLAB/glcSolve use deterministic search algorithms; therefore, the percentage of runs is not an applicable metric.

From the six solvers, TOMLAB/glcSolve solved our test problems quickest for both the centered and off-centered problem at 3rd and 37th function evaluations, respectively. In the best case, GPyOpt identified the optimal in under half of the evaluations needed by TOMLAB/glcSolve for the off-centered problem in 14 evaluations; however, the wall time necessary for GPyOpt to run for such a simple test problem was significantly longer than TOMLAB/glcSolve. MIDACO was second in identifying the optimum with the fewest function evaluations for the centered test problem with a median value of 256 evaluations and third for the off-centered problem with a median of 176 functional evaluations. Its hybrid search scheme allowed for quick identifications of potentially promising regions, then refining those regions for better solutions. Unfortunately, due to MIDACO's heavy reliance on random values for the identification of promising regions, it performed worse compared to the other solvers on reliably solving the problem, only solving 76 % of the cases for both test problems.

Table 3.11. Percentage of instances in which a solver found the global optimum, the evaluation number in which the solver found the global optimum, and the total invalid solutions the solver generated for A) the centered test problem and B) the 3 variable off-centered test problem

A) Centered Test Problem

	Percentage of runs optimum located	Evaluation # optimum identified at	Median number of invalid evaluations
DFL	N/A	935 N/A	0
NOMAD	98 %	317 (26, 673)	0
GPyOpt	78 %	410 (36, 997)	0
TOMLAB/glcSolve	100 %	3 N/A	0
MIDACO	76 %	256 (32, 904)	0
RBFOpt	78 %	433 (70, 966)	237

B) 3 Variable Off-centered Test Problem

	Percentage of runs optimum located	Evaluation # optimum identified at	Average number of invalid evaluations
DFL	N/A	44 N/A	0
NOMAD	98 %	189 (42, 806)	0
GPyOpt	94 %	312 (14, 982)	0
TOMLAB/glcSolve	100 %	37 N/A	0
MIDACO	76 %	176 (61, 895)	0
RBFOpt	86 %	456 (49, 1000)	256

Behind TOMLAB/glcSolve and DFL, NOMAD was the next most reliable of the solvers for the problems, finding the optimum 98 % of the time for both test cases. NOMAD and DFL, being local search solvers, require a multi-start approach to reach the allotted computational budget. These two solvers' performance relies heavily on the starting point, where DFL identified the optimum at evaluation 935 for the centered test problem and 44 for the off-centered problem. The performance of GPyOpt also depends on the starting points. Despite this dependence, GPyOpt

located the optimum for 78% runs for the centered test problem and 98% for the off-centered problem. The performance of RBFOpt was similar for both test problems, with a slight increase in the percentage of runs, from 78% to 86%, the optimum was located for the centered problem and off-centered problem. Due to its inability to handle constraints and categorical variables explicitly, RBFOpt required 433 function evaluations (the second-highest) for the centered and 456 evaluations (the highest) for the off-centered problem.

For this 4-variable combinatorial test problem, six solvers reliably identified the optimum solution within the computational budget. TOMLAB/glcSolve was found to be the quickest in the identification of the optimum for a combinatorial problem with a large number of symmetric solutions and numerous local minima. The solvers MIDACO and NOMAD were the next best two solvers for this problem. GPyOpt may not solve the problem quickly, but it is a top solver for reliably identifying the optimum solution. Finally, DFL and RBFOpt were able to solve the problem to optimality, but they experienced difficulties due to the symmetry of the problem. Furthermore, RBFOpt returned 256 infeasible solutions due to its inability to handle constraints, while the other five solver had none. Overall, TOMLAB/glcSolve performed the best out of the solvers used and is recommended for solving combinatorial optimization problems with numerous symmetric solutions and multiple local minima for a small number of variables. The TOMLAB suite, however, is a commercial solver package; if an open-source solver is preferred, the NOMAD solver is recommended. If the computational budget allows, the GPyOpt solver also shows promise for these types of problems. However, if the problem has a high number of variables, the solvers that should be considered are the open-source solvers NOMAD and DFL. For these, applying a multi-start approach with a space-filling sampling method would likely yield the most reliable results.

3.2.2 Evaluation of Selected DFO Solvers on CRC Screening Problem

3.2.2.1 CRC Screening Problem Description

The aim of the CRC screening problem is to identify a screening strategy for a defined population that maximizes the health benefits, quality-adjusted life years (QALY), of screening for CRC while minimizing the costs associated with CRC and screening. The decision variables are the screening test type, x_1, screening period, x_2, screening start age, x_3, and screening end age, x_4. The objective (Eq. 3.27) is to minimize the expected value of the difference between the total cost incurred due to CRC with screening and without screening, incorporating the QALY gained due to screening as a reward for gaining quality years of life for the simulated cohort. The expectation is approximated through the averaging of the total cohort costs calculated for each of the $|N|$ simulation iterations.

$$min\ \mathbb{E}\left(c^{Tot}(x_1, x_2, x_3, x_4)\right) \approx \frac{1}{|N|}\sum_{n=1}^{|N|}\left(\sum_{i=1}^{|I|} c_{i,n}^{Scr} - c_{i,n}^{Nat} - 100{,}000y'_{i,n}\right) \quad 3.27$$

$$s.t \quad c_{i,n}^{j} = c_{i,n}^{s,j} + c_{i,n}^{c,j} \qquad n \in N, \mathrm{i} \in I, j \in \{Scr, Nat\} \quad 3.28$$

$$y'_{i,n} = y_{i,n}^{Scr} - y_{i,n}^{Nat} \qquad n \in N, \mathrm{i} \in I \quad 3.29$$

$$c_{i,n}^{s,j} = f(x_1, x_2, x_3, x_4) \qquad n \in N, \mathrm{i} \in I, j \in \{Scr, Nat\} \quad 3.30$$

$$c_{i,n}^{s,j} = g(x_1, x_2, x_3, x_4) \qquad n \in N, \mathrm{i} \in I, j \in \{Scr, Nat\} \quad 3.31$$

$$y_{i,n}^{j} = h(x_1, x_2, x_3, x_4) \qquad n \in N, \mathrm{i} \in I, j \in \{Scr, Nat\} \quad 3.32$$

$$x_3 \leq x_4 \quad 3.33$$

$$x_1 \in \{\text{Colonoscopy, Sigmoidoscopy, FIT}\} \quad 3.34$$

$$x_2 \in \{1,2,\dots,40\} \quad 3.35$$

$$x_3 \in \{20,21,\dots,65\} \quad 3.36$$

$$x_4 \in \{50,51,\dots,85\} \quad 3.37$$

In Eq. 3.27, $c_{i,n}^{j}$ is the total incurred cost due to screening and treatment of CRC of individual i within the cohort, set I, for simulation iteration n, where the superscript j can either be Nat, denoting the lifetime where no screening strategy is implemented, or Scr, denoting the alternate lifetime in which a screening strategy is implemented. The variable $y'_{i,n}$ is the QALY gained due to screening for individual $i \in I = \{0,1,\dots,|I|\}$ of the simulated cohort for the n^{th} iteration of the microsimulation. For each decision variable set, the total incurred cost and QALY gained are calculated using the microsimulation, explained in Section 3.1. In Eq.3.28, $c_{i,n}^{j}$ is defined as the sum of screening costs, $c_{i,n}^{s,j}$, and costs due to cancer treatment, $c_{i,n}^{c,j}$ over the individual i's lifetime j. Eq. 3.29 calculates $y'_{i,n}$using QALY lived by individual i for lifetime j in simulation iteration n. The value of the QALY lived for the individual are based upon the health state of the individual pertaining to the progression of CRC. The values of $c_{i,n}^{s,j}$, $c_{i,n}^{c,j}$, and $y_{i,n}^{j}$ are determined by the microsimulation, and these relationships are represented by Eqns. 3.30 - 3.32. Equation 3.33 is a logical sequencing constraint ensuring that the starting age, x_3, for a screening strategy is less than or equal to the ending age, x_4. The domain for each decision variable are presented in Eqns.3.34 - 3.37.

3.2.2.2 Experimental Setup for the CRC Screening Problem

For each solver, the termination criterion was set to a maximum of 350 microsimulation evaluations, which allowed exploration of the search space within a reasonable computational time. The same screening strategy, a colonoscopy every 10 years starting at age 50 and ending at age 85, was used as the starting point for each solver, except for TOMLAB/glcSolve, which does not allow a choice for the starting point. For those solvers that initialized a random set of starting points, e.g., DEAP-GA, RBFOpt, and GPyOpt, the starting strategy was seeded within the initial set of points. We utilized a multi-start approach for solvers which converged before 350 evaluations from the specified starting point (e.g., DFL, and NOMAD) using a sequence of randomly pre-generated starting points. The solvers TOMLAB/glcSolve and TOMLAB/glcDirect perform identically when the problem is modeled using a categorical decision variable, as such, only results from TOMLAB/glcSolve are included in the analysis. The default parameter settings were used for the solvers as much as possible. When necessary, the parameters were set as detailed in Section 3.2.1.2 . The expected value in the objective function was estimated using 400 iterations of the CRC microsimulation, with each iteration having an identical random seed across all solvers. For each microsimulation evaluation, the cohort population was set to 1,000,000 males. A cohort of a single-sex was chosen to provide a screening strategy based on similar risk factors. A previous study (Meester et al., 2018) revealed that grouping cohorts by similar risk factors potentially allows for a more effective screening strategy than using a screening strategy for a generalized population. The entire cohort was initialized at birth, and each individual was tracked until age 100 or death, whichever condition occurs first. The compliance to screening was assumed to be 100 %, and the only screening tests performed to be the tests defined by the input screening strategy.

3.2.2.3 Results of the DFO Solver Evaluation for CRC Screening Problem

Table 3.12 displays the identified best screening strategy and its objective value, as well as the function evaluation at which the solver identified the solution for each of the tested DFO solvers. The ending age column shows a range rather than a single value, due to the symmetry present within the problem creating identical screening strategies within the range. It should be noted that the global optimum for this problem is unknown, and determining this value would require the evaluation of 19,644 unique screening strategies, or approximately 750 days' worth of computational time. As such, we will be referring to the minimum identified value as the optimum value, even though we cannot guarantee it is the global solution.

Table 3.12. The best-identified solutions, the evaluation number at which the DFO solver identified the solution and the number of invalid evaluations

Solver	Objective Value (Billion $)	Screening Test	Screening Period	Starting Age	Ending Age	Evaluation Number Identified	Invalid Evaluations
DFL	-10.29	Colonoscopy	11	38	82-85	115	0
NOMAD	-10.29	Colonoscopy	11	38	82-85	207	2
DEAP-GA	-10.13	Colonoscopy	11	28	83-85	40	50
GPyOpt	-10.11	Colonoscopy	13	37	76-85	229	1
SimAnneal	-10.12	Colonoscopy	13	39	78-85	194	50
glcSolve	-10.29	Colonoscopy	11	38	82-85	106	0
CMA-ES	-10.19	Colonoscopy	9	39	75-83	27	66
MIDACO	-10.29	Colonoscopy	11	38	82-85	112	2
RBFOpt	-10.06	Colonoscopy	10	28	78-85	346	107

The set of optimal solutions for the CRC screening problem was found to be a colonoscopy every 11 years, starting at the age of 38 with the ending age ranging from 82 to 85. The value of the objective function at the optimum was -10.29 billion dollars. The four solvers that identified the optimum were DFL, NOMAD, TOMLAB/glcSolve, and MIDACO, with each identifying the

solution at the 115th, 209th, 106th, and 112th function evaluations, respectively. The remaining solvers identified solutions that were within 3 % of the optimum. RBFOpt ended with the largest absolute percent difference and the highest invalid evaluation count (i.e., infeasible solutions) at 2.24 % and 107, respectively. The solution was identified at evaluation 346 and had an objective value of -10.06 billion dollars, suggesting to begin screening with a colonoscopy at age 28, every 10 years, with an ending age ranging from 78 to 85. The second-best performing solver, CMA-ES, identified a solution with an objective value of -10.19 billion dollars, absolute percent difference of 1.0 %, despite its second-highest number of invalid evaluations, 66. The screening strategy suggested to start screening with a colonoscopy at age 39 and screening every 9 years, ending screening from ages 75 to 83. The remaining three solvers, DEAP-GA, SimAnneal, and GPyOpt, had objective values of -10.13, -10.12, and -10.11 billion dollars, respectively. The solutions were identified at evaluations 40, 194, and 229 for DEAP-GA, SimAnneal, and GPyOpt, respectively. SimAnneal and DEAP-GA both came in third for the total invalid number of evaluations at 50, where GPyOpt only had one invalid evaluation.

We compare the performance of the solvers using two criteria. The first one is the change in the absolute percent difference from the optimum, given in Eqn. 3.38, with an increase in the number of function evaluations.

$$\% \, diff = \left| \frac{y_i^* - y^{opt}}{y^{opt}} \right| \tag{3.38}$$

In Eq. 3.38, $\% \, diff$ is the relative change and y_i^* is the current best-identified solution by the solver for function evaluation i, and y^{opt} is the optimum value. The second criterion aims to assess and compare the quality of the solutions, i.e., the screening strategies, generated by the solvers. We utilize QALY gained, total CRC treatment cost, total screening cost, preclinical CRC

incidence rate as a function of age, clinical CRC incidence rate as a function of age, and CRC death rate as a function of age to compare the effectiveness of the identified screening strategies.

3.2.2.3.1 Comparison of Required Simulation Evaluations to Find the Best Solution

The improvement of the absolute percent difference from the optimum as the number of function evaluations increases is plotted in Figure 3.24 for all solvers. Both NOMAD and DFL quickly identified a solution with an absolute percent difference under 5 %, with DFL reducing the difference to 2.2 % on evaluation 10 and NOMAD reducing the difference to 4.8 % on evaluation 11. Both solvers were able to identify the optimum, DFL was the third-fastest solver identifying the optimum at evaluation 115, and NOMAD was the slowest solver at finding the optimal at evaluation 207. The next solvers to identify solutions with an absolute percent difference under 5 % were TOMLAB/glcSolve, CMA-ES, and DEAP-GA. TOMLAB/glcSolve reduced the difference to 2.4 % on evaluation 36, CMA-ES reduced the difference to 1.0 % on evaluation 27, and DEAP-GA identified a solution with a 1.5 % difference at evaluation 40. Despite the quick identification of these solutions, both CMA-ES and DEAP-GA were unable to further improve their solutions for the remainder of the allotted budget. Consistent with results in Section 3.2.1.3 and the findings by Ploskas *et al.* (2018), TOMLAB/glcSolve identified the optimum the fastest at evaluation 106. Both GPyOpt and RBFOpt required a fair number of evaluations to construct a model that is able to identify points that improve the solution, similar to the observations in Section 3.2.1.3 . GPyOpt and RBFOpt required 118 and 102 evaluations, respectively, to identify a solution within 5 % of the optimum. GPyOpt, however, was able to identify a solution within 5.7 % at evaluation 20, where RBFOpt located a solution within 10 % at evaluation 90. Neither RBFOpt nor GPyOpt converged to the optimum within the computational budget. MIDACO, third to last for the initial identification of a solution with a less than 5 % difference, found a solution

with a 4.5 % difference at evaluation 57. However, MIDACO solved the problem second out of all of the solvers at evaluation 112. SimAnneal performed the worst, breaking the 5 % difference threshold at evaluation 123.

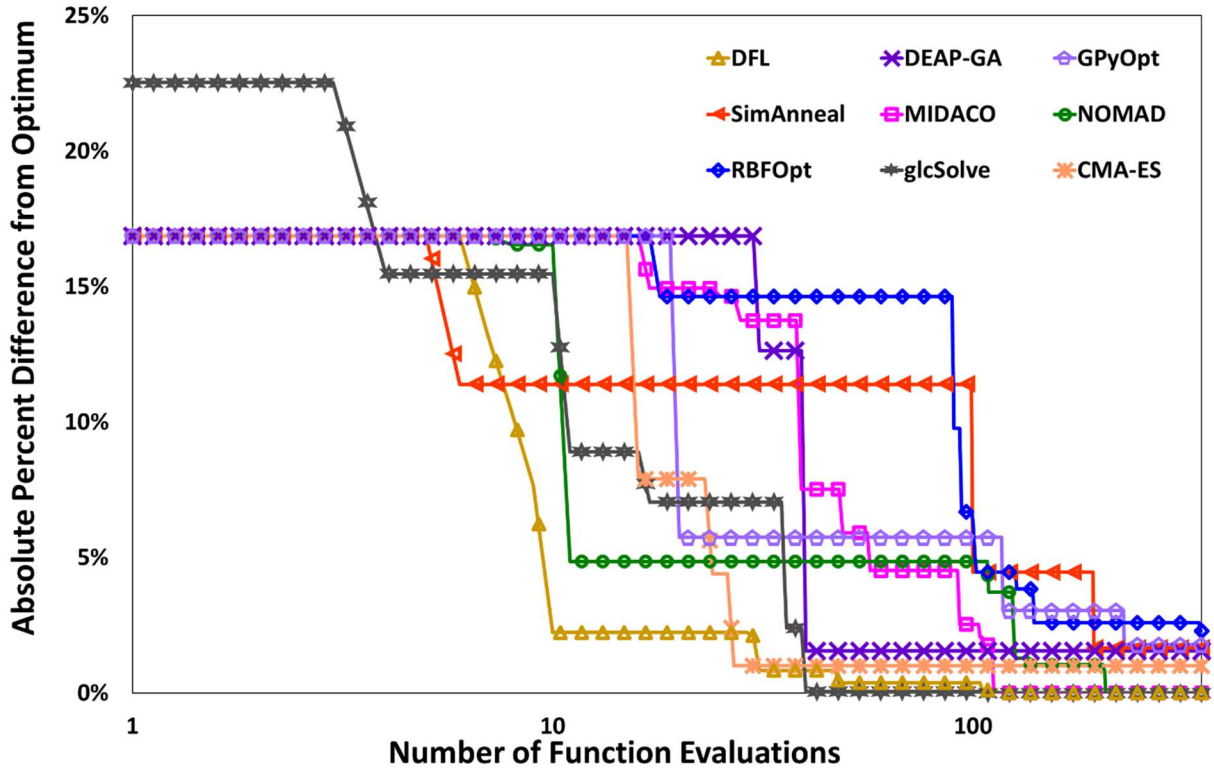

Figure 3.24 Change in absolute percent difference from the optimum with the number of function evaluations for all solvers for the CRC screening problem.

The computational budget was imposed to favor the solvers that efficiently explored the solution space, identifying regions that are promising for containing the optimum. Figure 3.25 (A) and (B) depicts the exploration of the solution space by the best, TOMLAB/glcSolve, and worst, DEAP-GA, solvers in terms of solution space explored. Plots showing the solution space explored by all solvers and each solver individually can be found in the supplementary material. The shape of the markers in Figure 3.25 correspond to the screening test type used, where a square denotes colonoscopy, a star denotes sigmoidoscopy, and a circle denotes FIT. The color of the markers is indicative of that point's objective value, with cooler colors being a lower objective value. Markers

that are pure black have an objective value of greater than $0/1,000,000 people. From Figure 3.25 (A), it can be seen that TOMLAB/glcSolve is very structured in its search path and is able to explore the solution space while exploiting the current best solution to sample solutions within that region. For DEAP-GA, Figure 3.25 (B), the random nature of the search strategy causes large gaps within the search space. Even when it found a good solution, DEAP- GA struggled to explore that area within the computational budget. Both SimAnneal and GPyOpt performed as two solvers that better explored the solution space. Unfortunately, due to its expansive exploration, GPyOpt was unable to identify the optimal solution within the budget. Similarly, SimAnneal was unable to identify the optimal solution as the process to generate new trial points is random, meaning it was unable to improve any identified promising solutions. One of the common traits shared by the solvers that found the optimum, MIDACO, NOMAD, DFL, and TOMLAB/glcSolve, was that they were all able to readily explore the region around a promising solution.

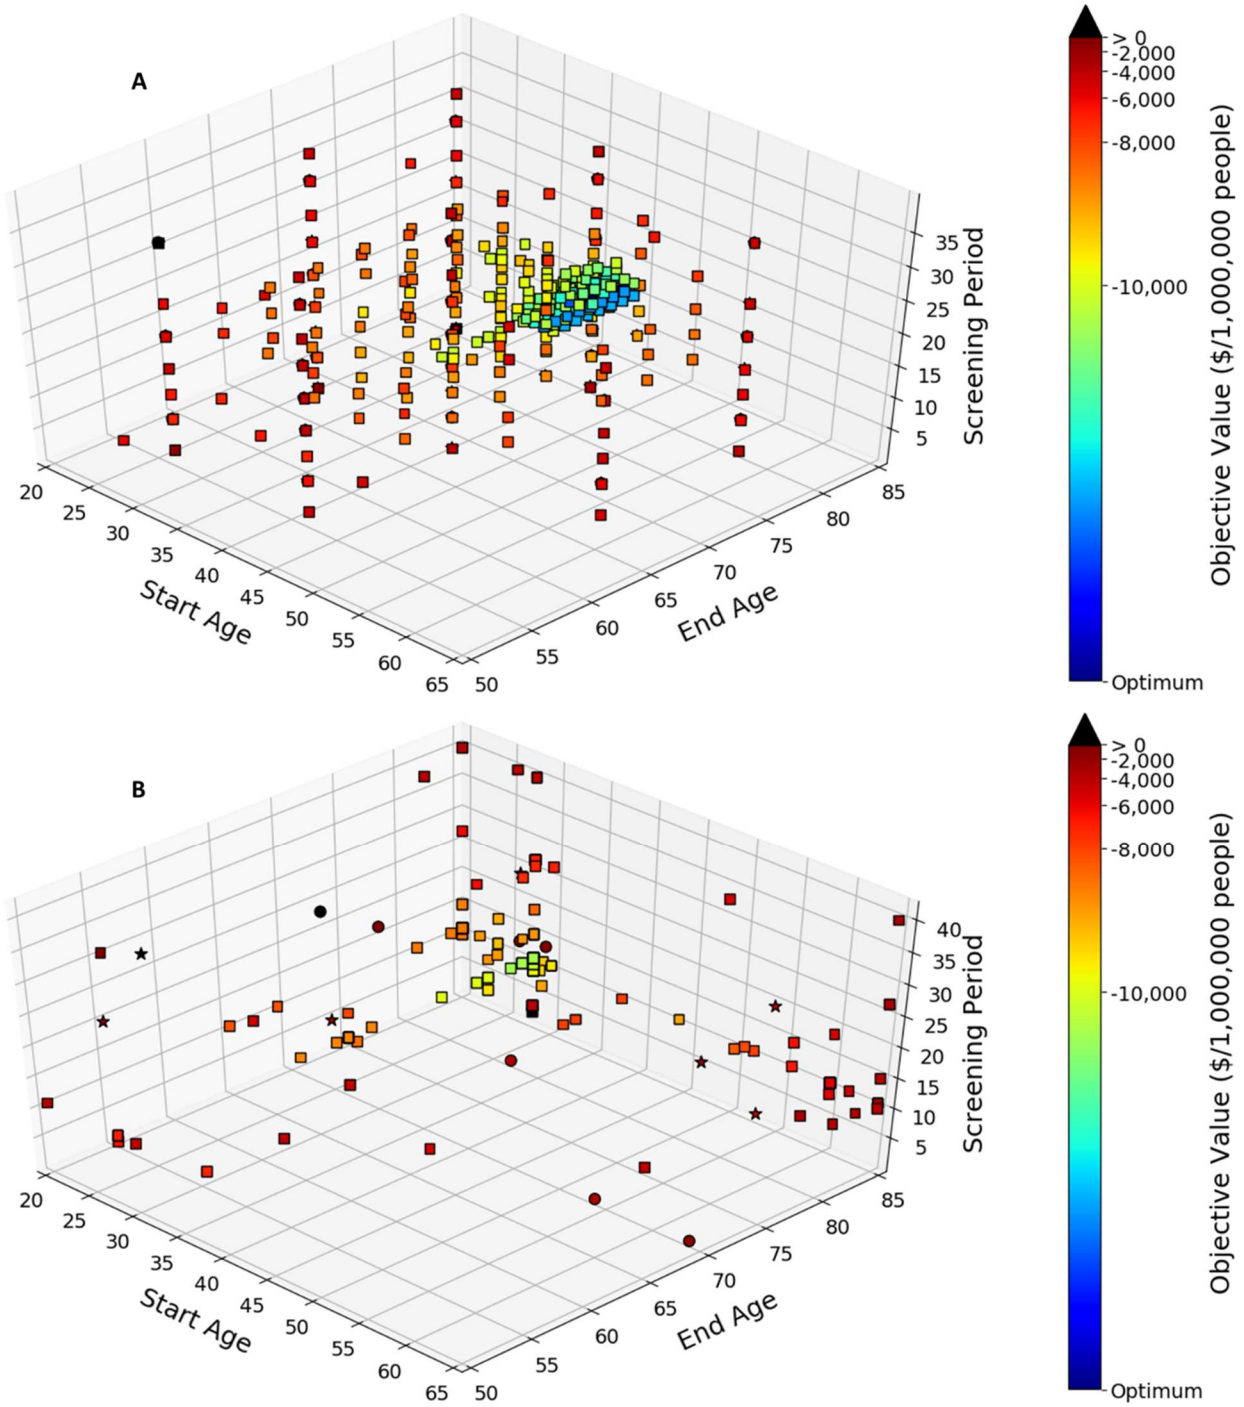

Figure 3.25 Exploration of the solution space with point markers corresponding to screening test type, square denotes colonoscopy, star denotes sigmoidoscopy, and circle denotes FIT, while color represents the objective function value. A) Most expansive solution space explored by the TOMLAB/glcSolve solver. B) Least explored solution space by the DEAP-GA solver.

Figures Figure 3.26 Figure 3.27 depict the progression of the absolute percent difference with the number of microsimulation evaluations for each of the starting points for DFL and

NOMAD, respectively. The first difference in the behavior of the two solvers can be seen just by the number of starting points needed by DFL to reach the 350 evaluations. DFL reaches convergence criterion for each run at a much lower number of evaluations in comparison to NOMAD when using default parameters. On average, the number of evaluation points per run for DFL was 44.3, while NOMAD had an average of 85.5 evaluations before termination due to reaching the minimum grid size. In Section 3.2.1, DFL was found to underperform due to its inability to escape local minima; however, for the CRC screening problem, DFL outperformed NOMAD. Given the four starting points that NOMAD used, NOMAD was only able to identify a better solution than DFL on a single run, Run 4. Of those four starting points, DFL located the optimum twice while NOMAD only located the optimum once. Furthermore, DFL was able to locate the optimum from four out of the eight of the starting points. A key search feature in NOMAD allows for any solution on the grid to be chosen as the next trial point. This allows the solver to escape local minima, which may also be the reason NOMAD was outperformed by DFL, as DFL stays within a local region of the incumbent solution, gradually moving along the search space improving the solution. This performance suggests that the CRC screening problem may be more well behaved than the test problems used in Section 3.2.1.

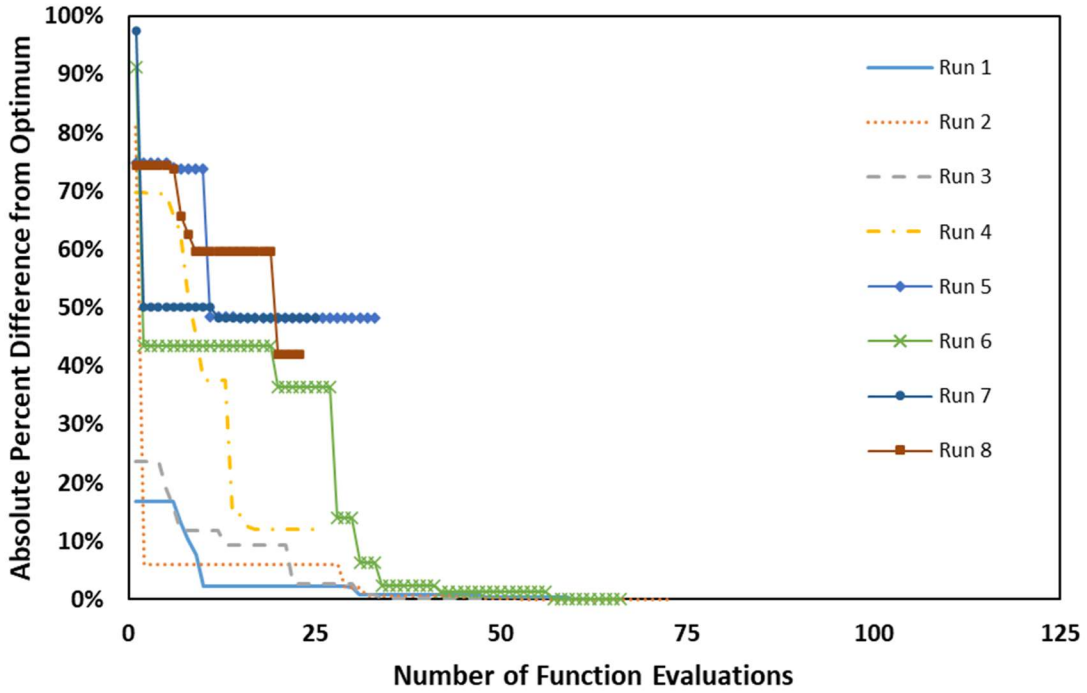

Figure 3.26 Change in absolute percent difference from the optimum with the number of function evaluations for the DFL solver using a different starting point for each run.

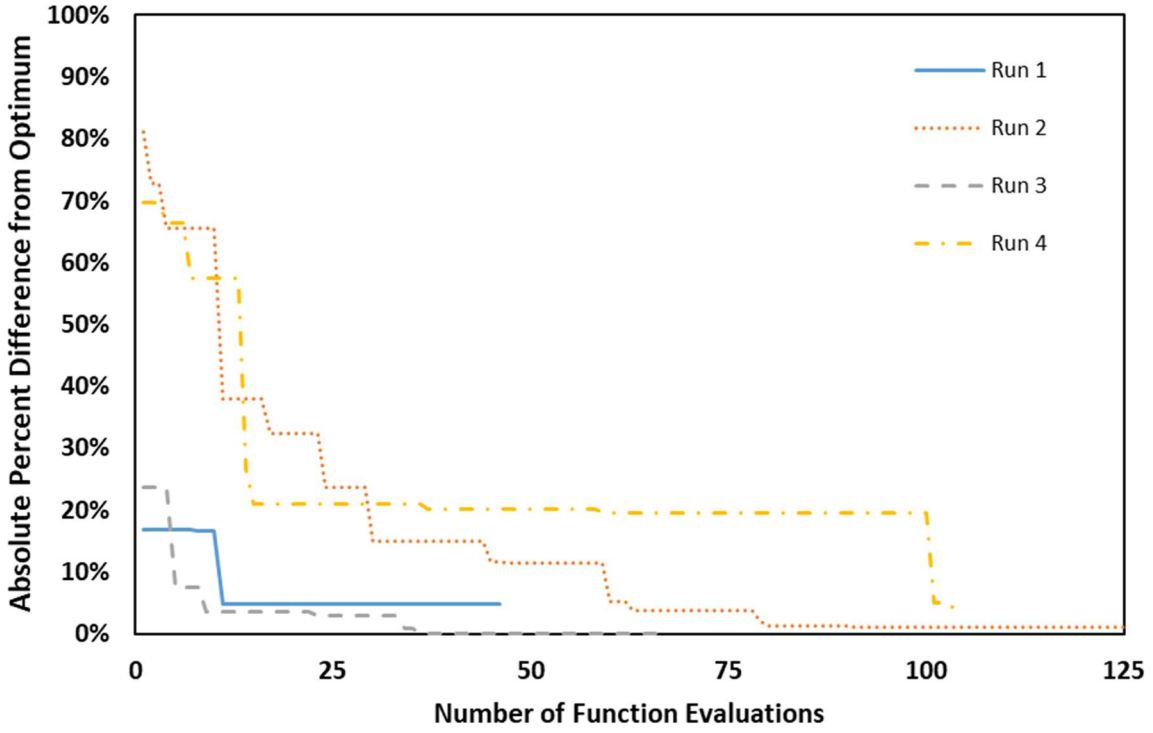

Figure 3.27 Change in absolute percent difference from the optimum with the number of function evaluations for the NOMAD solver using a different starting point for each run.

Figure 3.28 (A) and (B) present a 3-D surface representation of the objective function value scaled from $ to $ per 1,000,000 individuals. The color of the points represents the screening period, where a colder color represents less time in between screens and a warmer color represents more time in between screens. The shape of the points are similar to that in Figure Figure 3.25. Objective values greater than $0/1,000,000 people are all places within the same plane on the plot to discourage skewed bounds. Figure 3.28-(A) depicts the surface free of any scaling, while Figure 3.28-(B) applies log transformation to the objective function value (Z axis) to better visualize the global and local minima. The surfaces constructed from each solver's evaluations for both the unscaled and the log-scaled values can be found in the supplementary material. The shape of the CRC screening problem is smoother than our initial belief, with fewer than expected local minima observed. There is still a small degree of symmetric solutions within the problem. The surface demonstrates some discontinuities in the objective values, with more expected to be seen with an increase of evaluations. Overall, the CRC screening problem, though still challenging, is more well behaved than initially thought. The behavior of the problem allows for the solvers to find solutions within a 3 % difference from the optimal solution.

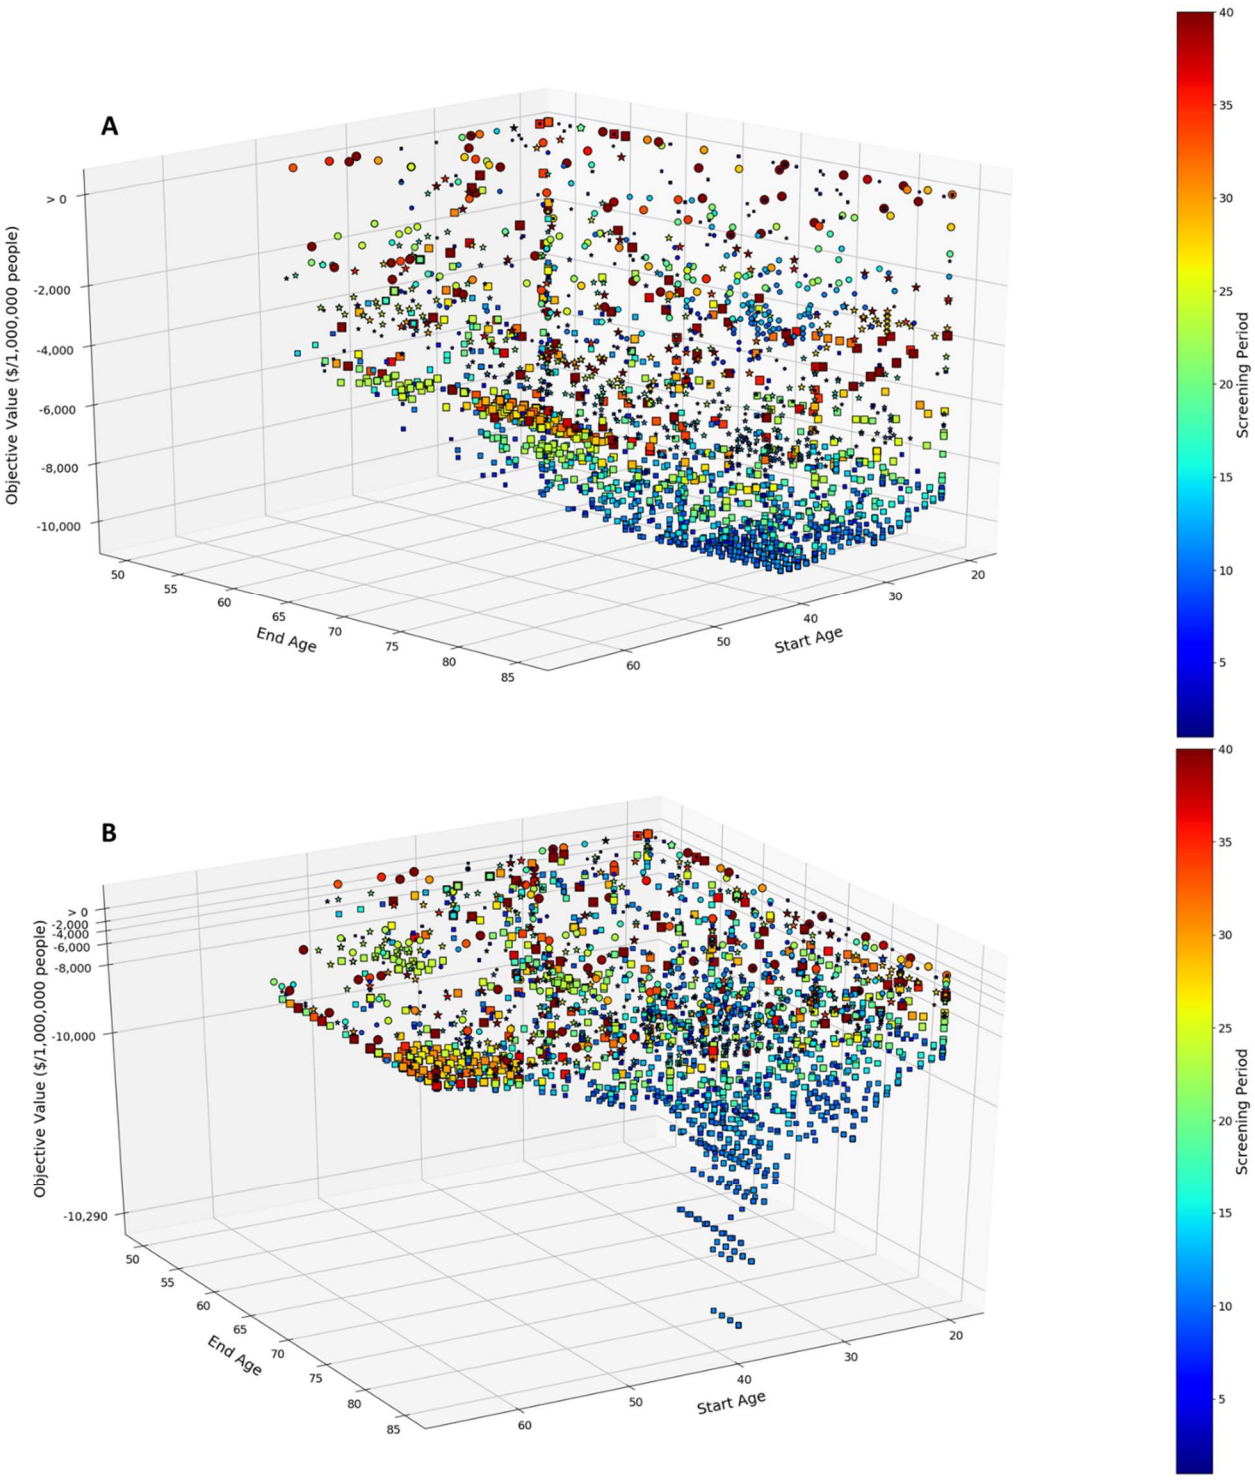

Figure 3.28 A) Objective function surface estimation using values obtained from all solver evaluations with color representing the screening frequency and the shape of the points representing the screening test type. **B)** Objective function surface estimation plotted with log scale to enhance the location of the global minima.

3.2.2.3.2 Comparison of the Solutions' Impact on CRC

Overall, each solver was able to obtain a solution that was within 2.5 % of the optimal objective value within the given computational budget. However, due to the competing terms of the objective function, i.e., reduction of total financial burden by minimizing cost versus reduction of the societal burden by maximizing the QALY gained, these solutions could vary greatly in the way that the overall reduction of burden is achieved. The further details of the top identified screening strategies are presented in Table 3.13., where strategy A is the optimal one identified by solvers DFL, NOMAD, MIDACO, and TOMLAB/glcSolve. Strategies B-F are identified by solvers CMA-ES, DEAP-GA, SimAnneal, GPyOpt, and RBFOpt, respectively, and strategy G is the current recommendation for the general population. In Table 3.13. are the values calculated by the simulation for the given screening strategy, including the QALY gained, the total cost of cancer care, and the total cost for the screening for a cohort of 1,000,000 men. All best-identified strategies (A-F) in Table 3.13. suggest a significantly earlier starting age for screening than the current recommendation, strategy G, while the screening frequencies are all relatively similar to strategy G. Unsurprisingly, the test of choice is a colonoscopy. Given perfect screening compliance, a colonoscopy is the most comprehensive and accurate test among available tests, even if it is the most expensive one [3]. Additionally, there is no need to perform a follow-up test on positive colonoscopy, unlike the other considered tests. The screening and treatment costs of the optimal strategy, strategy A, is around 5% higher than those of the strategy G, while the increase in the QALY gained is over 30%. For the QALY gained, the optimal strategy has the third highest value of all identified strategies. All of the strategies yield at least a 25 % increase in QALY gained compared to strategy G. The costs for treatment of cancer of all strategies is lower at least by 50 % than strategy G, with the optimal being only the fourth lowest cost. On the other hand, the

screening costs of the identified strategies using DFO solvers are nearly double that of the currently recommended strategy.

Table 3.13. Screening strategy description and outcome

Name	Test	Screening Period	Start Age	End Age	QALY gained	Cancer Cost (Million $)	Screening Cost (Million $)
A	Colonoscopy	11	38	82-85	100,854	314.48	1,307.69
B	Colonoscopy	9	39	75-83	100,569	306.50	1,417.88
C	Colonoscopy	11	28	83-85	104,960	279.85	1,938.42
D	Colonoscopy	13	39	78-85	97,720	413.48	1,093.45
E	Colonoscopy	13	37	76-85	98,475	415.36	1,177.44
F	Colonoscopy	10	28	78-85	105,343	265.53	2,067.44
G	Colonoscopy	10	50	80-85	76,498	871.15	660.08

The primary purpose of screening is to detect CRC at an early stage or to prevent CRC entirely by removing the adenoma before it becomes cancerous. To assess the benefit of a screening strategy solely on cost and QALY gained may not provide a complete view of its effects. The change in preclinical CRC rates, clinical CRC rates, and death rates due to CRC with each screening strategy are presented in Figure 3.29 (A), (B), and (C), respectively. These plots show the annual rates of the respective statistic as a function of the age of the cohort. Preclinical and clinical CRC rates represent the number of individuals who develop or are diagnosed with CRC at each year, respectively. In each plot, the black line represents the projected rate given no screening. Figure 3.29 reveals that each screening strategy drastically reduces the rates. The strategy with the minimum reduction of the rates, strategy E, still yields 70 %, 73 %, and 63 % reduction of the cumulative total cases for preclinical CRC, clinical CRC, and death due to CRC, respectively. Strategies F and C show the greatest cumulative reduction of all three measured rates, with C showing a greater reduction in the later ages of the cohort due to a later ending age screenings. Comparing strategy C to strategy A reveals a 4 % increase in QALY gained for the population. The percentage reduction of the cumulative total cases are very similar for both strategies A and

C, with Strategy C providing a 2 %, 1%, and 1% improvement from Strategy A in the reduction of the cumulative cases preclinical CRC, clinical CRC, and death due to CRC, respectively. However, screening and total costs of strategy C are 48 % higher than those of strategy A. The increase in cost is attributed to the extra screen recommended by the earlier starting age. Strategy A, as well as strategies B, D, and E, suggest starting screening in the age range from 37 to 39. This range corresponds to the ages in which the proportional development of adenomas in the population begins to slow, allowing for screening tests to detect these precursors prior to becoming cancerous, greatly benefiting the QALY gained. We plotted the total number of adenomas present in the simulated male cohort with no screening strategy implemented in Figure 3.30. The secondary axis in Figure 3.30 shows that the annual percent change in total adenomas reaches a change of about 8 % in the ages 37 to 39. The large decrease in total adenomas present in the population after the age of around 75 is consistent with the age at which much of the population begins to die via causes outside of CRC.

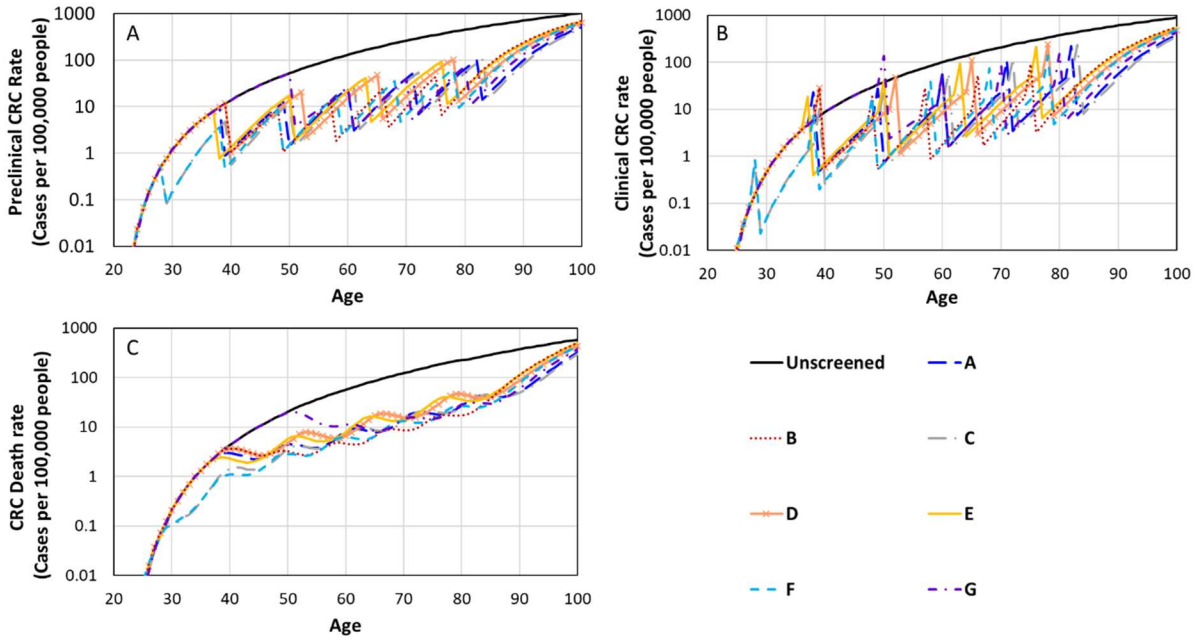

Figure 3.29 Male cohort simulation statistics from best-identified screening strategies A) Preclinical CRC incidence rate B) Clinical CRC incidence rate C) CRC Death rate

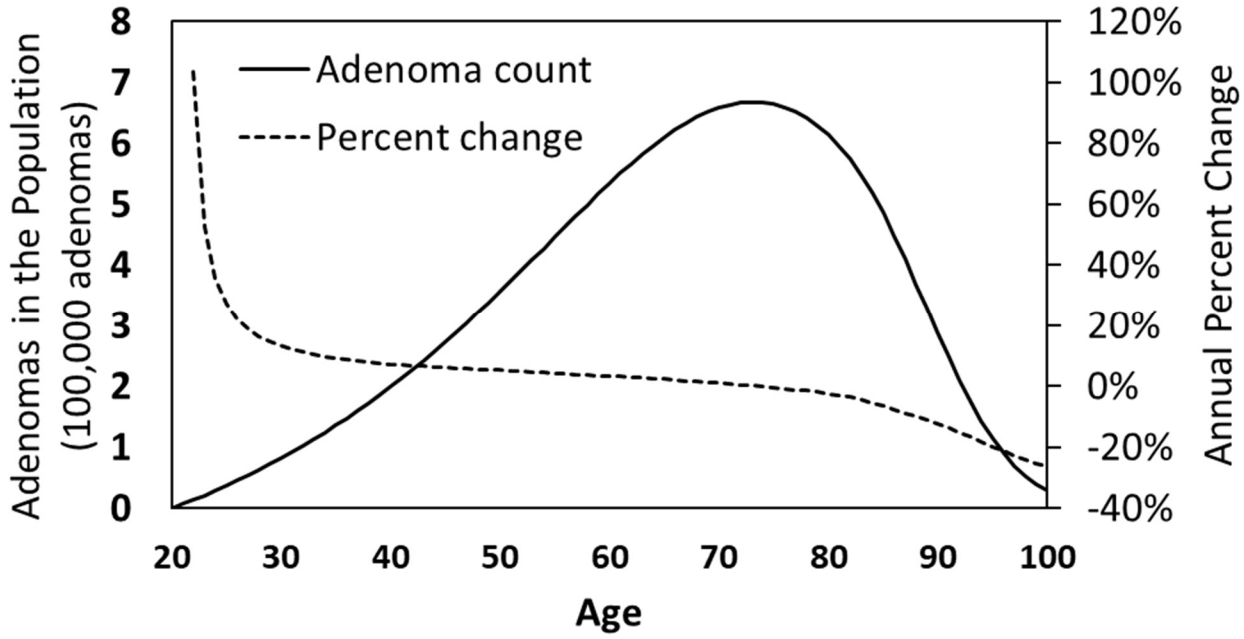

Figure 3.30 Total adenomas present in the unscreened male cohort as the cohort ages with the annual percent increase in total adenomas present in the population.

3.2.3 Study Findings

The results of this study were two fold. The first take away was a confirmation that the application of optimization methodology was able to determine CRC screening strategies that outperformed the currently recommended strategies. The best strategy identified from the DFO framework showed an increase in QALYG from screening by over 30 % while only raising the costs by around 5 %. The second take away was the DFO solvers that were best suited to solve the CRCSP. The two top performing solvers, TOMLAB/glcSolve and DFL, were able to identify a high quality solution in the fewest number of evaluations. GlcSolve was able to identify the region of a high quality solution quickly, though the convergence to the best solution would take longer. On the other hand, DFL was quick to converge on the locally optimal solution, but relied heavily on the starting point for the solver. If coupled together within the DFO, a combination of glcSolve and DFL would reliable be able to identify high quality, and locally optimal, solutions with a low number of function evaluations.

3.3 Simulation Sensitivity Study

Simulation optimization identifies solutions that exploit the structure and interactions of the simulation in order to generate the best desired output. If parameters or assumptions made within the simulation are modified, the optimal solution should also change.

3.3.1 Experimental set up

The DFO framework, determined in Section 3.2, exploits the CRC microsimulation to locate an optimum screening strategy considering the intrinsic simulation properties. The properties considered in the analysis are screening compliance implementation, the willingness-to-pay (WTP) ratio, and the risk ratio associated with the probability of an adenoma transitioning to a cancerous state. Additionally, the optimization problem is modified to include surveillance screening decisions variables to assess their impact. This study assesses the impact of different assumptions within microsimulations to understand how they affect the recommended policies as well as the optimal outcomes. Using the DFO framework, we study the effects of screening compliance modeling, the monetary value of a gain in lifespan, and adjusting the relative risk associated with a segment of the CRC progression pathway on the optimal screening strategy. We also assess the resulting screening and surveillance strategy when the surveillance decisions are included within the optimization problem. The results provide insight into how CRC microsimulation assumptions may result in selecting one screening policy over the others and highlights the importance of carefully assessing the assumptions in healthcare policy recommendations.

3.3.1.1 Optimization framework

The microsimulation was wrapped within a DFO framework to determine the optimal screening strategy for each of the tested parameters or assumptions. The framework uses two

solvers within the study. The first solver is glcSolve, from the commercial optimization suite TOMLAB[90]. It is used for investigating the compliance models, WTP ratio, and relative risk. GlcSolve utilizes the DIRECT (Dividing hyperRECTangles) algorithm[91], a global search DFO algorithm[80] that divides the search space into smaller hyper-rectangles, evaluating the quality of them by assessing the objective value of the midpoint. The second solver is an implementation of the DFL (Derivative-Free Linesearch) algorithm[81], a local search DFO algorithm[80] that aims to improve a solution by perturbing each decision variable individually by a given step size, slowly reducing the step size until a minimum step is reached. For integer variables this minimum step size is one. DFL was solely used to investigate the impact of incorporating surveillance decisions into the optimization problem. DFL was also used to refine the final solutions for each of the study to ensure local optimality of the solution.

The two solvers were chosen due to their ability to quickly and reliably converge to near optimal solutions for combinatorial problems[170]. In a previous study, it was found that for a low number of decision variables, glcSolve was able to identify the best solution for this optimization problem in under 500 evaluations. In the same study, it was found that DFL outperforms glcSolve as the number of decision variables increase when the computational budget is under 500 evaluations. The algorithmic parameters of glcSolve and DFL were kept at default values, and the maximum evaluations were set to 500 for the solvers, prior to solution refinement using DFL for each assumption or parameter value tested. During the refinement of the solution, DFL was allowed to run until the minimum step size was reached to ensure the solution is locally optimal.

3.3.1.1.1 Optimization problem

The optimization problem, Eqn. **3.39-3.43**, aims to identify the screening strategy that balances the gain of life from screening and the overall cost incurred due to CRC and screening

[170]. The decision variables are the screening test type (x_1), period (x_2), starting age (x_3), and ending age (x_4). The objective function, Eqn. **3.39**, minimizes the difference in total cost associated with CRC treatment and screening for a cohort for which the screening recommendation, c_n^{Scr}, is implemented and the same cohort without screening, c_n^{Nat}, while rewarding the gain in life from screening, y_n', with a WTP ratio per QALY gained ($/QALY), WTP, when screening was implemented. In Eqn. **3.39**, n is the n-th iteration of the microsimulation, and N is the total number of replications. Eqn. **3.40** calculates the change in QALY between the screened and unscreened cohort. The cost incurred and the QALY lived by the cohorts are calculated as part of the microsimulation (Eqn. **3.41**). Eqn. **3.42** enforces that the age at which screening begins is less than when it ends, and the bounds for each decision variable are defined in Eqn. **3.43**.

$$min\ \mathbb{E}\left(c^{tot}(x_1, x_2, x_3, x_4)\right) \approx \frac{1}{N}\sum_{n=1}^{N} c_n^{Scr} - c_n^{Nat} - (WTP)y_n' \qquad 3.39$$

$$s.t. \quad y_n' = y_n^{Scr} - y_n^{Nat} \qquad n \in N \qquad 3.40$$

$$c_n^j = f(x_1, x_2, x_3, x_4)$$
$$y_n^j = h(x_1, x_2, x_3, x_4) \qquad n \in N, \mathrm{j} \in \{Scr, Nat\} \qquad 3.41$$

$$x_3 \leq x_4 \qquad 3.42$$

$$x_1 \in \{Colonoscopy, Sigmoidoscopy, FIT\}$$
$$x_2 \in \{1,2, \dots ,40\}$$
$$x_3 \in \{20,21, \dots 65\}$$
$$x_4 \in \{50,51, \dots 85\} \qquad 3.43$$

3.3.1.2 Case Studies

The DFO algorithms used in the study assume that all inputs into the algorithm either have low noise or are deterministic. To decrease the noise of the output for the simulation, all analyses were carried out for a cohort of 1,000,000 males born in 1990, with the simulation replicated 400 times (N=400) to obtain a stable output value. The same simulated population was used to provide consistency amongst the evaluated screening strategies.

3.3.1.2.1 Screening compliance implementation

The study considered four compliance models for screening. The first is *perfect compliance*, where every individual is always compliant with the set screening strategy. The second utilizes a *flat compliance rate*, which states that an individual is either always compliant or never compliant based on the test type. The third is a *strict annual compliance* model, where an individual can always, sometimes, or never be compliant. If an individual is sometimes compliant, they can only be screened on the years set by the screening strategy. For example, if the strategy requires a screening every 10 years starting at age 50, the individual can only be screened ages 50, 60, …, etc. unless they undergo surveillance screening. The final compliance model is *loose annual compliance*, where an individual may always, sometimes, or never be compliant. In the loose annual compliance case, a sometimes-compliant individual may undergo a screening test any year after the defined time between screens has passed. For example, if an individual is to wait 10 years in between screens, a test may occur 10, 11, …, etc. years after the previous test unless the individual undergoes surveillance screening. The compliance rates employed are compiled in Table 3.14. As is evident, the compliance models increase in complexity, allowing better modeling of the compliance behavior and enabling the assessment of the effect on the screening recommendations.

Table 3.14 Adherence rates by test modality and screening model implementation.

	Flat Compliance [156]	Annual Compliance [19]		
Test modality	**Adherence rate**	**Always compliant**	**Sometimes Compliant**	**Annual adherence rate**
Colonoscopy	38 %	31 %	31 %	38 %
Sigmoidoscopy	38 %	31 %	31 %	38 %
FIT	67 %	40 %	29 %	50 %

3.3.1.2.2 Changes in willingness to pay ratio

The willingness to pay ratio (WTP) is ratio defined to be the maximum amount of money that a decision maker is willing to pay to gain a unit of utility, in this study QALYs. In our formulated optimization problem, the value of the WTP ratio is used to control the trade off between minimizing the overall cost and maximizing the gain in QALY by providing a monetary value to the change in QALYs due to screening. The CRC treatment and screening cost-effectiveness studies include WTPs in the range of \$20,000 – 100,000 per QALY[6,171]. We employed WTPs of \$20,000/QALY, \$50,000/QALY, \$75,000/QALY, and \$100,000/QALY using the loose annual compliance model. These values allowed us to investigate the range of WTPs for observing a general trend of its effect on the solution. The loose annual compliance model was chosen because it more realistically represents the individual compliance behavior.

3.3.1.2.3 Relative risk of adenoma transitioning to cancer

Modifying the relative risks (RRs) within the adenoma-carcinoma sequence enables assessing the impact an increase or decrease in the severity of the different parts of the CRC disease progression has on the optimal screening strategy for a cohort. For this study, we modified the RR associated with the probability of an adenoma transitioning into cancer. Three risk levels were considered, (1) base case, $RR = 1.0$, (2) 20 % increase, $RR = 1.2$, and (3) 20 % decrease, $RR =$

0.8. The value of +/- 20% of the RR was arbitrarily chosen to demonstrate the ability of the framework to capture the impact of changes in the simulation parameters. The perfect screening compliance model was chosen to better capture the effect of the risk modification by eliminating any variation in the screening strategy implementation across the cohort to show the maximum possible gain.

3.3.1.2.4 Incorporation of surveillance screening decisions

We modified the optimization problem, Eqns. **3.39** – **3.43**, to introduce three new decision variables, x_5, x_6, and x_7, for defining survelience screening periods for low, high, and special risk individuals (Table 1), respectively. The domain of these variables are defined in Eqn. **3.44**. We introduced constraint, Eqns. **3.45** – **3.47**, to reduce the solution space and enforce logical relationships for these variables. For example, Eqn. **3.45** states that the surveillance period for the low-risk category is less than or equal to the general screening period. The perfect screening compliance model was chosen to see the maximum benefit of incorporating the surveillance decisions.

$$x_5, x_6, x_7 \in \{1,2, \dots ,40\} \quad 3.44$$

$$x_5 \leq x_4 \quad 3.45$$

$$x_6 \leq x_5 \quad 3.46$$

$$x_7 \leq x_6 \quad 3.47$$

3.3.2 Results

3.3.2.1 The effect of screening compliance implementation

The results are summarized in Table 3.15. The total cost of CRC detection and treatment for the baseline, unscreened population is $1.855B. For perfect compliance, the best screening strategy identified by the DFO, Screening Strategy 1 (SS1), suggests a colonoscopy every 12 years, starting at age 38 and ending at 85. The objective function value is -$10.523B. Through screening, the total expected QALY gained is 102,328.20, with costs associated with screening $1.297B and with CRC treatment $0.268B, totaling to $1.565B, a reduction of 16 % compared to the unscreened population.

The strategies identified for the flat rate (SS2), strict (SS3), and loose (SS4) annual compliances include FITs. For SS2, an annual FIT is recommended between the ages of 31 to 85 yielding an objective function value of -$6.259B. Both SS3 and SS4 recommend an annual FIT from age 30 to 85 with an objective function value of -$6.091B. The total costs for SS2, SS3 and SS4 are all 14% lower than the cost of no screening. The expected QALY gained in SS2 and SS3/SS4 are 60,038.54 and 58,395.44 years, respectively. The results show that when assuming non-perfect compliance the best performing strategies favor the tests that the population is more likely to adhere to, even if the tests are less accurate and provide a reduction in the gain in QALY.

3.3.2.2 The effect of WTP

The screening strategies and associated metrics for evaluating the impact of changes in the WTP are summarized in Table 3.16. For all WTPs, the recommended strategy was an annual FIT with different start and end age ranges for screening: 30 to 85, 32 to 85, 33 to 84, and 38 to 85 for $100,000/QALY, $75,000/QALY, $50,000/QALY, and $20,000/QALY, respectively. The objective function values for these screening strategies were -$6.091B for a WTP of

$100,000/QALY, -$4.639B for $75,000/QALY, -$3.196B for $50,000/QALY, and -$1.501B for $20,000/QALY. The expected life years gained was 58,395.44 QALY for WTPs of $100,000/QALY, 57,898.35 QALY for WTP of $75,000/QALY, 57,572.27 QALY for WTP of $50,000/QALY, and 54,871.44 QALY for $20,000/QALY. The screening costs are $691MM, $642MM, $616MM, and $512MM for WTPs of $100,000/QALY, $75,000/QALY, $50,000/QALY, and $20,000/QALY. The total costs associated with screening and cancer care for the cohort was 14% less than the cost without a screening strategy for WTPs of $100,000/QALY, 16% less for $75,000/QALY, 17% less for $50,000/QALY, and 22% less for $20,000/QALY. As the reward for gaining QALY decreases the aggressiveness of the strategy also decreases, reducing the total screening tests given within a lifetime. This leads to a reduction in both screening costs and the gain in QALY.

3.3.2.3 The effect of relative risk of adenoma transitioning to cancer

Table 3.17 summarizes the results. The total costs without screening are $1.855B, $2.795B, and $1.032B for $RR = 1$, $RR = 1.2$, and $RR = 0.8$, respectively. In these instances, the total cost only includes costs incurred for CRC detection and care. The DFO framework identifies colonoscopy as the recommended screening modality every 12 years from age 38 to 85, 9 years from age 30 to 85, and 16 years from age 39 to 83 for $RR = 1$, $RR = 1.2$, and $RR = 0.8$, respectively. The corresponding objective function values and expected life-years gained (in the same order) are -$10.523B and 102,328.20 QALY, -$17.738B and 174,375.22 QALY, and -$4.898B and 50,728.59 QALY. Total costs incurred, including screening costs, decreased 16% to $1.565B for $RR = 1$, 11% to $2.495B for $RR = 1.2$, and increased 17% to $1.206B for $RR = 0.8$ when compared to their corresponding non-screened counterparts. Showing that as the average time for an adenoma to become cancerous decreases and the proportion of the adenomas

that become cancerous increases, the preferred initial screening age for the population shifts to earlier on in life, and the frequency at which test occur increases.

3.3.2.4 The effect of surveillance screening decisions

The optimal strategy determined was screening with a colonoscopy starting at age 38 and ending at age 84. The period between screens for the no risk category is 12 years, low risk is 7 years, high risk is 4 years, and special risk is 1 year. The objective function value is -$10.523B with the expected total life years gained of 102,313.78 QALY. The screening costs were $1.294B, with a CRC treatment cost of $270MM. Compared to SS1, the strategy with similar simulation assumptions but with a set surveillance strategy, there is a 0.01% reduction in QALY gained for the cohort, but also a 0.1% total cost reduction. This shows that the current recommendations for surveillance test to be near optimal and have little impact on the overall screening strategy.

3.3.3 Discussion

Of the three microsimulation assumptions studied, the optimum screening strategy identified was most sensitive to the changes in the RR, followed by the compliance modeling and WTP. Changes in the RR changed screening frequency and the screening starting and ending ages. Compliance modeling mainly impacted the screening modality. When WTP changed, the screening start age was affected.

Different screening compliance models introduce the human element into the analysis. Many studies utilize a perfect compliance model[2,21,172], obtaining a theoretical bound for the benefit due to screening strategies studied. Additionally, the optimum screening strategy located by an optimization algorithm identifies key aspects of the microsimulation that will significantly improve screening outcomes, such as critical ages to perform a screen, or the interval at which screening should occur when a perfect compliance model is utilized. However, incorporating more

intricate compliance models allows for a closer approximation of the real-world screening outcomes. Non-perfect screening compliance shifts the screening modality to the option that the most number of people are willing to take, which supports the claim that "Which test to use depends on the likelihood that you will get the test"[173]. For perfect compliance, colonoscopy, an accurate but invasive test with lower compliance rates[52,156], was recommended. In contrast, FIT, a less accurate and less invasive test with higher compliance rates[19,156], was recommended with other compliance models.

Table 3.15 Optimal screening strategies and statistics by the investigated compliance models.

Compliance Model	Screening Test	Screening Period	Start age	End age	Objective value ($ billion)	QALYG (years)	Screening Cost ($ billion)	CRC Cost ($ billion)	Total cost (No screening) ($ billion)
Perfect	Colonoscopy	12	38	85	-10.523	102,328.20	1.297	0.268	1.855
Flat	FIT	1	31	85	-6.259	60,038.54	0.721	0.879	1.855
Loose Annual	FIT	1	30	85	-6.091	58,395.44	0.691	0.913	1.855
Strict Annual	FIT	1	30	85	-6.091	58,395.44	0.691	0.913	1.855

Table 3.16 Optimal screening strategies and statistics for the investigated willingness to pay ratios using loose annual screening compliance.

Willingness to pay ($/QALY)	Screening Test	Screening Period	Start age	End age	Objective value ($ billion)	QALYG (years)	Screening Cost ($ billion)	CRC Cost ($ billion)	Total cost (No screening) ($ billion)
100,000	FIT	1	30	85	-6.091	58,395.44	0.691	0.913	1.855
75,000	FIT	1	32	85	-4.639	57,898.35	0.642	0.916	1.855
50,000	FIT	1	33	84	-3.196	57,572.27	0.616	0.922	1.855
20,000	FIT	1	38	85	-1.501	54,871.44	0.512	0.939	1.855

Table 3.17 Optimal screening strategies and statistics for the investigated relative risk values using perfect screening compliance.

Relative Risk	Screening Test	Screening Period	Start age	End age	Objective value ($ billion)	QALYG (years)	Screening Cost ($ billion)	CRC Cost ($ billion)	Total cost (No screening) ($ billion)
-20%	Colonoscopy	16	39	83	-4.898	50,728.59	1.042	0.164	1.031
Base	Colonoscopy	12	38	85	-10.523	102,328.20	1.297	0.268	1.855
+20%	Colonoscopy	9	30	85	-17.738	174,375.78	2.184	0.311	2.796

There is little difference in the screening strategies among the four WTPs for the loose annual compliance model. The only difference is the starting ages, later in life as the WTP decreases, which is not surprising because there is a lower likelihood of developing CRC earlier in life. However, the early developed CRC cases can significantly affect the QALY gained. For example, there is a 6% decrease in QALY gained when comparing the strategies for WTPs of $20,000/QALY and $100,000/QALY, where the only difference is a 6-year delay in the starting age. The similarities in the strategies suggest that the optimal strategy is not significantly affected by the WTP for the ranges considered.

The RR mainly impacts dwell time, the time it takes an adenoma to become cancerous, hence, affecting the proportion of the cohort that develops cancer. For $RR = 1$, the dwell time was 25.93 years on average, with a median of 24.08 years. For $RR = 1.2$, the dwell times dropped to an average of 24.32 years and a median of 22.35 years, while $RR = 0.8$ resulted in an average dwell time of 27.66 years with a median of 25.95 years. The 5th and 95th percentiles for the dwell time were 7.94 and 50.28, 8.82 and 52.38 and 7.17 and 48.24 years for $RR = 1$, $RR = 0.8$, and $RR = 1.2$, respectively. Furthermore, changes in RR significantly impact the QALY gained due to screening. A 20% increase in RR results in nearly a 50% increase in QALY gained, while a 20% decrease reduced QALY gained by 50%.

Screening frequency, starting age, and ending age vary in the recommended screening strategies for different RR. As the dwell time shortens, the recommended screening frequency increases. When the risk is higher ($RR = 1.2$), screening is recommended every 11 years, 45.2% of the average dwell time, in contrast to every 18 years, 65.1% of the average dwell time, for the lower risk ($RR = 0.8$). This change is likely due to the change in the proportion of adenomas that would become cancerous with RR. The optimum number of screens within an individual's lifetime

decreases as the risk decreases. Three screens are recommended for $RR = 0.8$, four for $RR = 1.0$, and five for $RR = 1.2$, suggesting that more intensive screening is optimal when the risk of developing cancer is higher or the dwell time is shorter.

There was very little change to the optimal screening strategy when considering the surveillance strategy as decision variables. The high-risk category had a slight increase in the screening period, from 3 to 4 years, while the special category had a reduction, from 2 to 1 year. This slight change improved the cost-effectiveness by 0.6% and improved the objective function value by 0.001%. These results suggest that the surveillance strategy used was near-optimal, given the assumption of perfect screening compliance.

3.4 **Conclusions**

This chapter outlines the creation and evaluation of a simulation based optimization approach for solving the CRCSP. The constructed simulation shows to approximate the original CRC-SPIN simulation allowing a good estimate for the progression of CRC within a population. The simulation is easily able to incorporate in-depth details of the screening process or progression of CRC. When incorporated into a DFO framework, the simulation can easily be used to generate high quality screening strategies that may out-perform current recommended strategies.

Of the solvers, the commercial solver TOMLAB/glcSolve performed the best overall, identifying the global optima for the test problems with four variables the quickest. The performance of TOMLAB/glcSolve did deteriorate as the number of variables in the problem increased. Although the commercial solver MIDACO was able to solve the test problems and the CRC screening problem, it suffered from a lack of consistency due to the random nature of the main search strategy when restricted to a relatively low number of function evaluations. The solver's inconsistency of finding a solution can be compensated for by how quickly the solver is

able to generate and evaluate new trial points. If the problem is computationally inexpensive, multiple runs of MIDACO can be performed for the same amount of time a single run of NOMAD, GPyOpt, or RBFOpt would take.

The next notable solver is the open-source solver NOMAD, which was the most reliable solver behind TOMLAB/glcSolve. It was able to find the global optima of both four variable problems within a reasonable number of function evaluations. Additionally, NOMAD was a top performer as the number of variables in the problem increased alongside DFL, while the remaining solvers' performances quickly deteriorated with the increase. The solvers that utilize a model-based search strategy, i.e., GPyOpt and RBFOpt, are outperformed by the direct solvers for problems whose black-box models are not expensive to calculate due to the large time needed by the solver for the determination of the next trial point. However, with a large number of allowed evaluations, these solvers prove to be effective. When the problem is computationally expensive, the time handicap model-based solvers have compared to the direct search solvers are less prevalent.

The application of a DFO framework with a CRC microsimulation model was used to determine the optimal screening strategy for CRC for a cohort of males. The identified strategies by the DFO framework, though more expensive than the currently recommended one, suggests that there may be a potential for improving the benefits of CRC screening using an optimization-based approach. The cost of cancer treatment was reduced by 64 %, with an increase in quality-adjusted life-years gained by 31 % when using the optimal strategy for a male cohort as identified by the DFO framework compared to the current general recommendation for CRC screening. Additionally, the optimal solution showed a 30 % reduction in total cases of CRC and death attributed to CRC when compared to current recommendations.

Microsimulation models provide valuable information for assessing healthcare policy options. Coupling these microsimulations with an optimization framework enables the complex interactions in the simulation, not seen by researchers, to be exploited for obtaining optimal healthcare policies using the microsimulation. Assumptions and parameters of the microsimulation can be modified to investigate the impact these changes have on the identified optimal policy. The changes in optimal policy can be used to identify aspects of the microsimulation that have the greatest impact on the solution. Those aspects can then be refined to provide better policies and more accurately estimates of the outcomes of the policies.

The differences in the optimal screening strategies demonstrated the impact of assumptions made in a microsimulation, with some causing a more significant impact than others. The results revealed that the changes in the disease progression have the most impact, suggesting that optimal screening strategies would differ when an individual's risk factors for CRC are considered. The RR modifications in this study were hypothetical. Performing a similar study while incorporating the results of clinical trials to assess the effects of CRC risk factors and their impacts on the disease progression could provide valuable information and pave the way for developing personalized CRC screening strategies.

CHAPTER 4

MATHEMATICAL PROGRAMMING

This chapter introduces the mathematical programming framework to solve the CRCSP. The problem is presented in two different formulations. The first formulation, Section 4.1, presents a two-stage stochastic programming approach to solve the CRCSP incorporating both exogenous and type I endogenous uncertainty into the problem. The second formulation, Section 4.2, presents a multi-stage stochastic programming (MSSP) approach for solving the CRCSP. This approach incorporates surrogate modeling to approximate a closed form expression for the benefits associated with screening an asymptomatic population for CRC.

4.1 Two-stage Stochastic Programming Approach to Solve CRCSP

We define a class of problems for scheduling state-identifying actions to detect system abnormalities with imperfect outcomes by maximizing system performance as an imperfect test scheduling problem (ITS). There are two main sources of uncertainty in ITS, the uncertainty of the system's state at any given time and the sensitivity of the action in detecting the system state. One application domain of the ITS is a healthcare related-problem, scheduling diagnostic tests to identify asymptomatic individuals with a disease, such as the CRCSP. Section 4.1.1 presents a general formulation of the ITS and then extends the model to investigate the CRCSP.

4.1.1 General Problem Formulation and Model Description

In ITS, a set of actions is scheduled to identify if a system is in an abnormal state that will decrease its utility. The actions impose a disutility on the system and are imperfect, making the identification of the true state of the system uncertain. The probability of observing the system's

true state and the system's disutility are both functions of the action taken and the system's current state. The system's progression through the states is known. However, the time spent in each state is unknown but follows a known distribution per stage. Suppose the system is detected to be in an abnormal state. In that case, a rectifying action is taken to return the system to its nominal state or to a state that maximizes the system utility. It is also assumed that if the system has a rectifying action taken on it, it no longer needs identification actions to be taken on it within the scheduling timeframe.

4.1.1.1 Generalized MINLP formulation (MINLP1)

We formulate the problem as a two-stage stochastic program with fixed recourse. We define set $J = \{1, \dots, j\}$ to be the set of all possible identification actions the scheduler can take, set $Z = \{1, \dots, z\}$ to be the set of all possible action numbers, set $T = \{1, \dots, t\}$ to be the timeframe considered in the scheduling of the actions, and set $S = \{1, \dots, s\}$ as all potential system states. The first stage decision variables, $X_{j,t,z} \in [0,1]$, identify if action j is taken at time point t as the z-th action on the system. In the problem, the decisions must be properly ordered (e.g., action z must occur after action $z-1$). This constraint is presented in Eqn. 4.1. Additionally, for each of the $|Z|$ actions, one action type must be chosen, as shown in Eqn. 4.2.

$$\sum_{t \in T} \sum_{j \in J} X_{j,t,z} = 1 \qquad \forall z \in Z \qquad .1$$

$$\sum_{j \in J} X_{j,t,z} \leq \sum_{j \in \boldsymbol{J}} \sum_{b=20}^{t} X_{j,t,z} \qquad \forall t \in T, \forall z \in Z \qquad 4.2$$

The system's state progression through the states is represented by the uncertain parameter $\theta_{s,t,i}$, where $i \in I$ is a finite set of scenarios from the full uncertainty space, $\mathcal{I}$, of all possible realizations of the system progression. The probability measure of $\boldsymbol{\theta}_i$, $\mathbb{P}_i$, is known for all potential realizations. The outcomes of the actions taken are represented by $\Xi_{z,\omega}$, where $\Xi_{z,\omega} = 1$ if action z detects the true state of the system for outcome ω and is zero otherwise. In this, $\omega \in \Omega$ is the set of all possible combinations of outcomes of actions taken on the system, such that $|\Omega| = 2^{|Z|}$. The probabilities of a given outcome, $\mathbb{P}_{s,j}(\Xi_{z,\omega})$, given the system's state, s, and the action taken, j, are known with certainty. However, the probability measure of the outcome of the actions is, $\mathbb{P}(\boldsymbol{\Xi}_\omega|\boldsymbol{X}, \boldsymbol{\theta}_i)$, unknown until actions are taken and depends on the timing of the actions taken. The expression for determining the probability measure of $\boldsymbol{\Xi}_\omega$ is presented in Eqn. 4.3.

$$\mathbb{P}(\boldsymbol{\Xi}_\omega|\boldsymbol{X}, \boldsymbol{\theta}_i) = \prod_{z\in Z} \sum_{j\in J} \sum_{s\in S} \sum_{t\in T} \mathbb{P}_{s,j}(\Xi_{z,\omega}) \theta_{s,t,i} X_{j,t,z} \tag{4.3}$$

The second stage variables, $Y_{s,t,i,\omega}$, are used to represent the updated stage progression of the system, where $Y_{s,t,i,\omega} = 1$ if the system is at state s at time point t for realization i given outcome ω of the actions. The values of the second stage variables are determined through system-dependent fixed rules. The general form of the fixed recourse action constraints can be found in Eqn. 4.4

$$Y_{s,t,i,\omega} = g(\boldsymbol{X}, \boldsymbol{\theta}_i, \boldsymbol{\Xi}_\omega) \qquad \forall s \in S, t \in T, i \in I, \omega \in \Omega \tag{4.4}$$

The deterministic equivalent MINLP, MINLP1, can be written as Eqns 4.5 – 4.8.

$MINLP1$:

$$max \sum_{i\in I} \mathbb{P}_i \sum_{\omega\in\Omega} \left(\prod_{z\in Z} \sum_{j\in J} \sum_{s\in S} \sum_{t\in T} \mathbb{P}_{s,j}(\Xi_{z,\omega}) \theta_{s,t,i} X_{j,t,z} \right) F(\boldsymbol{X}, \boldsymbol{Y}_{i,\omega}, \boldsymbol{\theta_i}, \omega) \quad .5$$

$$s.t. \quad 4.1,4.2,4.4 \quad .6$$

$$X_{j,t,z} \in [0,1] \qquad \forall j \in J, \forall t \in T, \forall z \in Z \quad .7$$

$$Y_{s,t,i,\omega} \in [0,1] \qquad \forall s \in S, t \in T, i \in I, \omega \in \Omega \quad .8$$

4.1.1.2 Linearization of the objective function

Due to the inclusion of type I endogenous uncertainty, the objective function is inherently nonlinear. As the number of actions taken within the time frame increases, so does the problem nonlinearity.

4.1.1.2.1 Linearization of binaries then continuous (MILP1)

To obtain the probability measure of the results of the first stage actions, Eqn. 4.3 can be reformulated using exact linearization[174] of the product of binary variables through the introduction of binary variables $\boldsymbol{X}^{L1}$, as shown in Eqns. 4.9 and 4.10. This new variable tracks the system's stage and the action type for each of the $|Z|$ actions taken in the first stage per realization of the system's state progression. The probability of the outcomes of the first stage actions can then be reconstructed, as seen in Eqn. 4.11, through the introduction of the new parameter, $\mathbb{P}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),\omega}$, where $\mathbb{P}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),\omega}$ is the probability that outcome ω occurs when

action j_1 is taken when the system is in state s_1 and so on for the remaining ($|Z| - 1$ action, system state) pairs.

$$X^{L1}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),i} \leq \sum_{t \in T} \theta_{s_z,t,i} X_{j_z,t,z} \qquad \forall z \in \boldsymbol{Z}, \forall (s_1,j_1), \dots, (s_{|Z|},j_{|Z|}) \in (S \times J)^{|Z|}, \forall i \in I \tag{4.9}$$

$$X^{L1}_{(s_1,\dots,s_{|Z|}),(j_1,\dots,j_{|Z|}),i} \geq \sum_{t \in T} \sum_{z' \in Z} \theta_{s_z,t,i} X_{j_z,t,z'} - (|Z| - 1) \qquad \forall z \in \boldsymbol{Z}, \forall (s_1,j_1), \dots, (s_{|Z|},j_{|Z|}) \in (S \times J)^{|Z|}, \forall i \in I \tag{4.10}$$

$$\mathbb{P}(\boldsymbol{\Xi}_\omega | \boldsymbol{X}, \boldsymbol{\theta}_i) = \sum_{\substack{(s_1,j_1),\dots,\\ (s_{|Z|},j_{|Z|}) \in (S \times J)^{|Z|}}} \mathbb{P}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),\omega} \, X^{L1}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),i} \tag{4.11}$$

We can then linearize the term $X^{L1}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),i} F(\boldsymbol{X}, \boldsymbol{Y}_{i,\omega}, \boldsymbol{\theta_i}, \omega)$ in Eq. 4.5 through the introduction of a new continuous variable, $U^{L1}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),i,\omega}$, by the exact linearization[175] of the product of a binary and continuous variable as seen in Eqns. 4.12 – 4.15. Parameters M_i^U and M_i^L are the upper and lower bounds for the big M constraint of state realization i for the system that can be calculated as the best- and worst-case utility for each realization of the state progression.

$$U^{L1}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),i,\omega} \leq M_i^U X^{L1}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),i} \qquad \forall (s_1,j_1), \dots, (s_{|Z|},j_{|Z|}) \in (S \times J)^{|Z|}, \forall i \in I, \forall \omega \in \Omega \tag{4.12}$$

$$U^{L1}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),i,\omega} \geq M_i^L X^{L1}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),i} \qquad \forall (s_1,j_1), \dots, (s_{|Z|},j_{|Z|}) \in (S \times J)^{|Z|}, \forall i \in I, \forall \omega \in \Omega \tag{4.13}$$

$$U^{L1}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),i,\omega} \leq F(\boldsymbol{X}, \boldsymbol{Y}_{i,\omega}, \boldsymbol{\theta_i}, \omega) - M_i^L \left(1 - X^{L1}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),i}\right) \qquad \forall (s_1,j_1), \dots, (s_{|Z|},j_{|Z|}) \in (S \times J)^{|Z|}, \forall i \in I, \forall \omega \in \Omega \tag{4.14}$$

$$U^{L1}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),i,\omega} \geq F(\boldsymbol{X}, \boldsymbol{Y}_{i,\omega}, \boldsymbol{\theta_i}, \omega) - M_i^U \left(1 - X^{L1}_{(s_1,j_1),\dots,(s_{|Z|},j_{|Z|}),i}\right) \tag{4.15}$$

$$\forall (s_1, j_1), \dots, (s_{|Z|}, j_{|Z|}) \in (S \times J)^{|Z|}, \forall i \in I, \forall \omega \in \Omega$$

Using the above linearization approach, we remove the nonlinear terms within the objective function, Eqn. 4.5, of MINLP1 at the expense of introducing a number of new variables and constraints. This results in the following equivalent MILP model, Eqns. 4.16 – 4.17, referred to as MILP1.

MILP1:

$$max \sum_{i \in I} \mathbb{P}_i \sum_{\omega \in \Omega} \sum_{\substack{(s_1, j_1), \dots, \\ (s_{|Z|}, j_{|Z|}) \in (S \times J)^{|Z|}}} \mathbb{P}_{(s_1, j_1), \dots, (s_{|Z|}, j_{|Z|}), \omega} U^{L1}_{(s_1, j_1), \dots, (s_{|Z|}, j_{|Z|}), i, \omega} \quad 4.16$$

$$s.t. \; 4.1, \mathbf{4.2}, 4.4, 4.7 - 4.15 \quad 4.17$$

4.1.1.2.2 Probability chain linearization (MILP2)

The probability chain approach[128] uses a recursive calculation to incrementally linearize the probability-weighted utility function through the introduction of the two sets of continuous variables, $X^{L2}_{s,j,z,i,\omega}$ and $U^{L2}_{z,i,\omega}$. In this approach, we use $U^{L2}_{z,i,\omega}$ as an accounting variable that calculates the probability-weighted utility for outcome ω for system realization i up action z. Here, $U^{L2}_{|Z|,i,\omega}$ is the full probability-weighted utility for outcome ω for system realization i given the first stage decisions, while $X^{L2}_{s,j,z,i,\omega}$ is used to properly calculate the probability weighting for each action taken on the system. Equations 4.18 – 4.22 show the constraints used to calculate the first term of the probability-weighted utility, $U^{L2}_{1,i,\omega}$. The direct linearization of a binary and continuous variable is given in Eqns. 4.18 – 4.21, where the probability-weighted utility is calculated for each action type and system state combination. However, the value of $X^{L2}_{s,j,1,i,\omega}$

remains zero unless the action type and system state match the type of action chosen by action one, and the system state matches the system's state when action one is taken. This results in Eqn. 4.22, properly calculating the probability-weighted utility due to action one.

$$X^{L2}_{s,j,1,i,\omega} \leq M^U_i \sum_{t\in T} \theta_{s,t,i} X_{j,t,1} \qquad \forall s \in S, \forall j \in J, \forall i \in I, \forall \omega \in \Omega \quad 4.18$$

$$X^{L2}_{s,j,1,i,\omega} \geq M^L_i \sum_{t\in T} \theta_{s,t,i} X_{j,t,1} \qquad \forall s \in S, \forall j \in J, \forall i \in I, \forall \omega \in \Omega \quad 4.19$$

$$X^{L2}_{s,j,1,i,\omega} \leq \mathbb{P}_{s,j}(\Xi_{1,\omega}) F(\boldsymbol{X}, \boldsymbol{Y}_{i,\omega}, \boldsymbol{\theta_i}, \omega) + M^U_i \left(1 - \sum_{t\in T} \theta_{s,t,i} X_{j,t,1}\right) \quad \forall s \in S, \forall j \in J, \forall i \in I, \forall \omega \in \Omega \quad 4.20$$

$$X^{L2}_{s,j,1,i,\omega} \geq \mathbb{P}_{s,j}(\Xi_{1,\omega}) F(\boldsymbol{X}, \boldsymbol{Y}_{i,\omega}, \boldsymbol{\theta_i}, \omega) + M^L_i \left(1 - \sum_{t\in T} \theta_{s,t,i} X_{j,t,1}\right) \quad \forall s \in S, \forall j \in J, \forall i \in I, \forall \omega \in \Omega \quad 4.21$$

$$U^{L2}_{1,i,\omega} = \sum_{s\in S} \sum_{j\in J} X^{L2}_{s,j,1,i,\omega} \qquad \forall i \in I, \forall \omega \in \Omega \quad 4.22$$

We can then formulate similar constraints for actions two up to $|Z|$, building on the probability-weighted utility from the previous action. These constraints are presented in Eqns. 4.23 – 4.27.

$$X^{L2}_{s,j,z,i,\omega} \leq M^U_i \sum_{t\in T} \theta_{s,t,i} X_{j,t,z} \qquad \forall z \in Z, \forall s \in S, \forall j \in J, \forall i \in I, \forall \omega \in \Omega \quad 4.23$$

$$X^{L2}_{s,j,z,i,\omega} \geq M^L_i \sum_{t\in T} \theta_{s,t,i} X_{j,t,z} \qquad \forall z \in Z, \forall s \in S, \forall j \in J, \forall i \in I, \forall \omega \in \Omega \quad 4.24$$

$$X^{L2}_{s,j,z,i,\omega} \leq \mathbb{P}_{s,j}(\Xi_{z,\omega}) U^{L2}_{z-1,i,\omega} + M^U_i \left(1 - \sum_{t\in T} \theta_{s,t,i} X_{j,t,z}\right) \quad \forall z \in Z, \forall s \in S, \forall j \in J, \forall i \in I, \forall \omega \in \Omega \quad 4.25$$

$$X^{L2}_{s,j,z,i,\omega} \geq \mathbb{P}_{s,j}(\Xi_{z,\omega}) U^{L2}_{z-1,i,\omega} + M^L_i \left(1 - \sum_{t\in T} \theta_{s,t,i} X_{j,t,z}\right) \quad \forall z \in Z, \forall s \in S, \forall j \in J, \forall i \in I, \forall \omega \in \Omega \quad 4.26$$

$$U^{L2}_{z,i,\omega} = \sum_{s\in S} \sum_{j\in J} X^{L2}_{s,j,z,i,\omega} \qquad \forall z \in Z, \forall i \in I, \forall \omega \in \Omega \quad 4.27$$

With the linearization methodology presented in Eqns. 4.18 – 4.27, we can once again replace the nonlinearities in the original objective function, Eqn. 4.5, and obtain a larger but linear problem. The resulting equivalent MILP is presented in Eqn. 4.28 and 4.29 and referred to as MILP2.

MILP2:

$$max \sum_{i \in I} \mathbb{P}_i \sum_{\omega \in \Omega} U^{L2}_{|Z|,i,\omega} \quad 4.28$$

$$s.t.\ 4.5, 4.20, 4.21, 4.23 - 4.27 \quad 4.29$$

4.1.1.3 Extension to Colorectal Cancer Screening Problem (CRCSP)

The objective of the CRCSP is to maximize the expected quality-adjusted life-years (QALY) gained by scheduling screening actions within an individual's lifetime. Quality-adjusted life years is a weighted measure of years an individual has lived based on the health state an individual is in for a given year of life[19]. The gain in QALY is calculated by comparing the QALY of an individual without any screening actions to the QALY if screening actions were taken within their life. The weights associated with the years lived are tied to the progression of CRC. As CRC progresses, the quality-of-life weight decreases.

For the TSSP of CRCSP, we assume that if a non-healthy state, adenoma, or undiagnosed CRC is detected for any individual, that person can no longer be screened within their lifetime. For the recourse actions taken after a screening action is performed, nothing occurs if a healthy state is detected. We assume that an individual would return to a healthy state if an adenoma is detected, removing any future possibility of developing CRC. If the screen detects CRC, that individual transitions to a clinically detected CRC state corresponding to the stage of CRC

detected. The survival years with CRC are then modified to fit the survival probabilities for the new stage of CRC and the age at which CRC was diagnosed[8]. Additionally, the only screening action considered is screening with a colonoscopy to simplify the model formation and because it is the most accurate of the tests for CRC identification[3].

4.1.1.3.1 Problem Nomenclature

The first stage decisions, $X_{j,t,z}$, are the screening test type, j, an individual is prescribed to take at age t and which test number it would be, z. The constraints fixing the recourse actions are the clinical actions that would be taken on an individual upon the screening test results for CRC. The second stage binary variables, $Y_{s,t,i,\omega}$, represent the resulting lifetime for individual i given screening outcome ω by identifying if they are at health state s at age t or not. The set of states, $s \in S = [1,2,\dots,11]$, represent the different health states associated with the progression of CRC given in Figure 2.1 and are further described in Table 4.1. The age of the individual is represented by the set $t \in T = [20,21,\dots,100]$, and the set $i \in I \subset \mathcal{I}$ represents the different individuals that are prescribed the screening recommendation in a subset of the full population. The set $\omega \in \Omega$ represents the different possible screening outcomes. The available screening actions are constrained to one option, a colonoscopy, i.e., $j = J = [1]$, where 1 represents a colonoscopy. In addition to these sets, we also include sets: $s_d \in S_d = [1,2,..,6]$ as the subset of health states that are detectable by screening tests, $d \in D = [1,2,3,4]$ as the set of death sources from CRC that correspond to the CRC stage an individual is clinically detected at, and $k \in K = [1,2,3,4,5]$ as the set of age brackets that are used for CRC survival post clinical detection. The age brackets used to define set K are as follows: $1 = [20,50), 2 = [50,60), 3 = [60,70), 4 = [70,80), 5 = [80,100)$.

The parameters in SCRCSP are c_s, d_{j,s_d}, $\kappa_{k,t}$, $\Xi_{z,\omega}$, $\mathbb{P}_i$, and $\mathbb{P}_{s,j}(\Xi_{z,\omega})$. Parameters c_s and d_{j,s_d} provide the quality-of-life measure for spending a year within state s and a quality-of-life

disutility for screening test j detecting health state s_d, respectively. The parameter $\kappa_{k,t}$ takes the value of one if age t falls in age bracket k and is zero otherwise. The parameter $\Xi_{z,\omega}$ is used to represent the outcome of screening action z for the outcome realization ω, where $\Xi_{z,\omega} = 1$ if screening action z identifies non-healthy states for outcome ω and is zero otherwise. $\mathbb{P}_i(\boldsymbol{\theta}_i)$ and $\mathbb{P}_{s,j}(\Xi_{z,\omega})$ are used to define the probability measure of the realization of individual i and the probability measure of outcome $\Xi_{z,\omega}$ occurring when action j is being used to identify state s, respectively. The values for parameters c_s, d_{j,s_d}, and $\mathbb{P}_{s,j}(\Xi_{z,\omega})$ are presented in Table 4.1.

Table 4.1. Model parameters and health state definitions.

State	c_s	$d_{j=1,s_d}$	$\mathbb{P}_{s,j=1}(\mathbf{0})$	$\mathbb{P}_{s,j=1}(\mathbf{1})$
1 – Healthy	1.00	0.0055	1.0	0.0
2 – Adenoma	1.00	0.0055	0.1	0.9
3 – Undiagnosed stage 1 CRC	0.88	0.0055	0.01	0.99
4 – Undiagnosed stage 2 CRC	0.88	0.0055	0.01	0.99
5 – Undiagnosed stage 3 CRC	0.88	0.0055	0.01	0.99
6 – Undiagnosed stage 4 CRC	0.88	0.0055	0.01	0.99
7 – Clinical stage 1 CRC	0.80	-	1.0	0.0
8 – Clinical stage 2 CRC	0.75	-	1.0	0.0
9 – Clinical stage 2 CRC	0.60	-	1.0	0.0
10 – Clinical stage 2 CRC	0.30	-	1.0	0.0
11 – Deceased	0.00	-	1.0	0.0

4.1.1.3.2 Model Uncertainty

The uncertain parameters for this problem are $\theta_{s,t,i}$, $\tau_{t,i}^{Nat}$, and $\tau_{d,t,k,i}^{CRC}$. Parameter $\theta_{s,t,i}$ takes a value of one if individual i is at health state s at age t and zero otherwise, $\tau_{t,i}^{Nat}$ takes a value of one if individual i would have passed away from sources outside of CRC at or before age t and zero otherwise, and $\tau_{d,t,k,i}^{CRC}$ takes a value of one if individual i would have passed away from stage d CRC, that was detected in age bracket k at or before age t and zero otherwise. The values for $\boldsymbol{\theta}_i$

and $\boldsymbol{\tau}_i^{CRC}$ are estimated using data generated from a microsimulation model to construct an empirical distribution of potential health-state progression outcomes, while the values for $\boldsymbol{\tau}_i^{Nat}$ are estimated from an empirical distribution constructed from the 2000 USA lifetables[176]. The microsimulation model used implements the first iteration of the Colorectal Cancer Simulated Population model for Incidence and Natural history (CRC-SPIN)[8,177].

CRC-SPIN is a continuous-time simulation that models the progression of CRC within a defined population. Within each individual, multiple adenomas are allowed to develop over their lifetime, with each adenoma having the potential to develop into CRC. Each potential case of CRC within an individual has an associated stage, time until detected through symptomatic means (Sojourn time), and survival time, with the CRC that would be symptomatically detected first determining the overall survival of CRC. The empirical distribution for generating $\boldsymbol{\theta}$ is constructed by recording the following in a 5-dimensional (5-D) vector for every individual simulated within the model: 1) the age at which the first adenoma occurs, 2) the age at which the first undiagnosed CRC occurs, 3) the age at which the first CRC that becomes detected, 4) the CRC stage that is detected, and 5) the age at which the individual would die from that case of CRC. Each 5-D vector is treated as a discrete realization of a possible progression of CRC within an individual. If the progression of CRC is an unrealistic realization, such as the age of the first adenoma is an age older than 100, we default that realization to a default realization of (101,101,101,0,101) where CRC does not occur within the individual. We then construct an empirical distribution around the discrete realizations of CRC progression, simplifying the 5-D multivariate distribution into a single variate distribution. We employed 400 replications of the microsimulation to generate the vectors, with each replication simulating the lifetime of 1,000,000 males born in 1990.

The survival time for CRC within the simulation model used depends on the age and CRC stage at detection. To generate the values for $\boldsymbol{\tau}_i^{CRC}$, we use the values 3, 4, and 5 of the empirical distribution used to construct $\boldsymbol{\theta}_i$ to back-calculate the survival probability for individual i from the simulation. We apply the same survival probability to generate the ages at which the individual would pass away due to CRC given an earlier detection of the disease for both age of detection and stage of detection for CRC. The ages are generated using the same survival model used within the microsimulation.

4.1.1.3.3 Formulation

The utility function of CRCSP measures the QALY gained due to the first stage decisions. This is expressed in Eqn. 4.30, where c_s and d_{j,s_d} are constants that provide the quality of life measure for spending a year within state s and a quality of life disutility for screening test j detecting health state s_d, respectively. Variable $\nu_{s_d,z,k,i,\omega}$ is a binary variable used to identify which state, s_d, screening test z detected for individual i in age bracket k for screening outcome ω.

$$F\left(\boldsymbol{X}, \boldsymbol{Y}_{i,\omega}, \boldsymbol{\theta_i}, \omega\right) = \sum_{t\in T}\sum_{s\in S} c_s(Y_{s,t,i,\omega} - \theta_{s,t,i}) - \sum_{z\in Z}\sum_{j\in J}\sum_{s_d\in s_d} d_{j,s_d}\sum_{k\in K}\nu_{s_d,z,k,i,\omega} \qquad 4.30$$

Equations 4.31 to 4.35 define $\nu_{s_d,z,k,i,\omega}$. Equations 4.31 and 4.33 define if the screening actionhas identified the initial healthy state ($s = 1$), for the first screening test and the second screening test and onward respectively. In Eqn. 4.31, the first term of the expression determines if the screening outcome, $\boldsymbol{\Xi}$, correctly identifies that individual i is in a healthy state, while the

second term is used for when the screening test incorrectly identifies the individual as healthy due to a flase negative. . Eqns. 4.32 and 4.34 sets $\nu_{s_d,z,k,i,\omega}$ to one if the screening action z identifies a non-healthy state s_d for the first action ($z = 1$) and any subsequent actions ($z \geq 1$), respectively. Finally, Eqn. 4.35 ensures that only one state be identified per screening action. The subtraction of the term $\sum_{s'_d=2}^{6}\sum_{z'=1}^{z}\sum_{k'\in k}\nu_{s'_d,z',k',i,\omega}$ in Eqns. 4.33 to 4.35 disallow further screening actions to be taken on individual i that had a previous action identify a non-healthy state.

$$\nu_{1,1,\mathrm{k},\mathrm{I},\omega} = \Xi_{1,\omega}\sum_{j\in J}\sum_{t\in T}\kappa_{k,t}\theta_{1,t,\mathrm{i}}X_{j,t,1} + \left(1-\Xi_{1,\omega}\right)\left[\sum_{j\in J}\sum_{t\in T}\kappa_{k,t}X_{j,t,1} - \sum_{j\in J}\sum_{t\in T}\sum_{s'=7}^{11}\kappa_{k,t}\theta_{s',t,\mathrm{i}}X_{j,t,1}\right] \quad \forall k\in K, \forall i\in I, \forall\omega\in\Omega \tag{4.31}$$

$$\nu_{\mathrm{s_d},1,\mathrm{k},\mathrm{I},\omega} = \Xi_{1,\omega}\sum_{j\in J}\sum_{t\in T}\kappa_{k,t}\theta_{\mathrm{s_d},t,\mathrm{i}}X_{j,t,1} \quad \forall s_d\in S_d|s_d\geq 2, \forall k\in K, \forall i\in I, \forall\omega\in\Omega \tag{4.32}$$

$$\nu_{1,z,k,i,\omega} \geq \Xi_{z,\omega}\sum_{j\in J}\sum_{t\in T}\kappa_{k,t}\theta_{1,t,i}X_{j,t,1} - \sum_{s'_d=2}^{6}\sum_{z'=1}^{z}\sum_{k'\in k}\nu_{s'_d,z',k',i,\omega} + \left(1-\Xi_{z,\omega}\right)\left[\sum_{j\in J}\sum_{t\in T}\kappa_{k,t}X_{j,t,1} - \sum_{j\in J}\sum_{t\in T}\sum_{s'=7}^{11}\kappa_{k,t}\theta_{s',t,i}X_{j,t,z}\right] \quad \forall z\in Z|z\geq 2, \forall k\in K, \forall i\in I, \forall\omega\in\Omega \tag{4.33}$$

$$\nu_{s_d,z,k,i,\omega} \geq \Xi_{z,\omega}\sum_{j\in J}\sum_{t\in T}\kappa_{k,t}\theta_{s_d,t,1}X_{j,t,1} - \sum_{s'_d=2}^{6}\sum_{z'=1}^{z}\sum_{k'\in K}\nu_{s'_d,z',k',i,\omega} \quad \forall s_d\in S_d|s_d\geq 2, \forall k\in K, \forall z\in Z|z\geq 2, \forall i\in I, \forall\omega\in\Omega \tag{4.34}$$

$$\sum_{s_d\in S_d}\sum_{k\in K}\nu_{s_d,z,k,i,\omega} \leq 1 - \sum_{s'_d=2}^{6}\sum_{z'=1}^{z}\sum_{k'\in K}\nu_{s'_d,z',k',i,\omega} \quad \forall z\in Z|z\geq 2, \forall i\in I, \forall\omega\in\Omega \tag{4.35}$$

We introduce a binary variable $\psi_{z,t,i,\omega}$ that tracks if individual i has screening test z identify a non-healthy state at or before age t given outcome ω, the value of this variable is set

through Eqns. 4.36 to 4.41. Equations 4.36 and 4.37 set the value of $\boldsymbol{\psi}$ to one for the screening test $z = 1$ if the first screening action detects a non-healthy state at age $t = 20$ and $t \geq 21$, respectively. Equations 4.38 and 4.39 force the value of $\boldsymbol{\psi}$ to be one for the screening test $z \geq 2$ if a previous test has detected a nonhealthy state at age $t = 20$ and $t \geq 21$, respectively. Equations 4.40 and 4.41 provide a check to evaluate if screening test $z \geq 2$ has identified a nonhealthy state regardless of previous test results.

$$\psi_{1,20,i,\omega} = \Xi_{1,\omega} \sum_{s=2}^{6} \theta_{s,20,i} \sum_{j \in J} X_{j,20,1} \qquad \forall i \in I, \forall \omega \in \Omega \qquad 4.36$$

$$\psi_{1,t,i,\omega} = \Xi_{1,\omega} \left(\sum_{s=2}^{6} \theta_{s,t,i} \sum_{j \in J} X_{j,t,1} + \psi_{1,t-1,i,\omega} \right) \qquad \forall t \in T | t \geq 21, \forall i \in I, \forall \omega \in \Omega \qquad 4.37$$

$$\psi_{z,20,i,\omega} \geq \Xi_{z,\omega} \left(\sum_{s=2}^{6} \theta_{s,20,i} \sum_{j \in J} X_{j,20,z} - \sum_{z'=1}^{z} \psi_{z',20,i,\omega} \right) \qquad \forall z \in Z | z \geq 2, \forall i \in I, \forall \omega \in \Omega \qquad 4.38$$

$$\psi_{z,t,i,\omega} \geq \Xi_{1,\omega} \left(\sum_{s=2}^{6} \theta_{s,t,i} \sum_{j \in J} X_{j,t,z} + \psi_{z,t-1,i,\omega} - \sum_{z'=1}^{z} \psi_{z',t,i,\omega} \right) \qquad \forall t \in T | t \geq 21, \forall z \in Z | z \geq 2, \forall i \in I, \forall \omega \in \Omega \qquad 4.39$$

$$\psi_{z,20,i,\omega} \leq \Xi_{z,\omega} \sum_{s=2}^{6} \theta_{s,20,i} \sum_{j \in J} X_{j,20,z} \qquad \forall z \in Z | z \geq 21, \forall i \in I, \forall \omega \in \Omega \qquad 4.40$$

$$\psi_{z,t,i,\omega} \leq \Xi_{z,\omega} \left(\sum_{s=2}^{6} \theta_{s,t,i} \sum_{j \in J} X_{j,t,z} + \psi_{z,t-1,i,\omega} \right) \qquad \forall t \in T | t \geq 21, \forall z \in Z | z \geq 2, \forall i \in I, \forall \omega \in \Omega \qquad 4.41$$

Variable $\xi_{t,i,\omega}$ is introduced to track if individual i is deceased at t for outcome ω. Equation 4.42 enforces the age of death due to causes outside of CRC if any screening actions detect an unhealthy state. Equation 4.43 enforces the age of death due to CRC based upon the stage of CRC detected by the screening actions.

$$\xi_{t,i,\omega} \geq \tau_{t,i}^{Nat} \sum_{z\in Z}\sum_{k\in K}\sum_{s=2}^{6} \nu_{s,z,k,i,\omega} + \left(1 - \sum_{z\in Z}\sum_{k\in K}\sum_{s=2}^{6} \nu_{s,z,k,i,\omega}\right)\theta_{11,t,i} \qquad \forall t \in T, \forall s \in S, \forall \omega \in \Omega \qquad 4.42$$

$$\xi_{t,i,\omega} \geq \tau_{d,t,k,i}^{CRC} \sum_{z\in Z} \nu_{s,z,k,i,\omega} \qquad (s_d, d) \in [(3,1),(4,2),(5,3),(6,4)], \; \forall t \in T, \forall k \in K, \forall i \in I, \forall \omega \in \Omega \qquad 4.43$$

With the tracking variables $\boldsymbol{\nu}$, $\boldsymbol{\psi}$, and $\boldsymbol{\xi}$ defined for key events for CRC prevention or early detection, we now define the intersection of key events that need to be enforced to properly define the resulting lifetime of an individual due to the screening actions taken. We define these intersections directly by introducing two new sets of variables, $\boldsymbol{\eta}$ and $\boldsymbol{\phi}$, to directly model the product of binary variables linearly. It should also be noted that these logical AND events could be implemented through disjunctive programming.

The first set of variables, $\eta_{s_d,t,i,\omega}$, is the product of $\nu_{s_d,z,k,i,\omega}$ and $\psi_{t,i,\omega}$ and is used to track if health state s_d was detected for individual i at or before age t given screening outcome ω. The direct linearization[174] of the product of the binary variables $\nu_{s_d,z,k,i,\omega}$ and $\psi_{t,i,\omega}$ is given in Eqns. 4.44 to 4.48. The next variable, $\phi_{s_d,t,i,\omega}$, is used to track if individual i has health state s_d detected at or before age t given screening outcome ω and if that individual has passed away at or before that age t. This variable is represented by the product of $\eta_{s_d,t,i,\omega}$ and $\xi_{t,i,\omega}$, with the constraints resulting from direct linearization[174] of this product presented in Eqns. 4.49 to 4.51.

$$\eta_{s_d,t,i,\omega} \leq \sum_{k\in K}\sum_{z\in Z} \nu_{s_d,z,k,i,\omega} \qquad \forall s_d \in S_d | s_d \geq 2, \forall t \in T, \forall i \in I, \forall \omega \in \Omega \qquad 4.44$$

$$\eta_{1,t,i,\omega} \leq \sum_{k\in K} \nu_{1,z,k,i,\omega} \qquad \forall t \in T, \forall z \in Z, \forall i \in I, \forall \omega \in \Omega \qquad 4.45$$

$$\eta_{s_d,t,i,\omega} \leq \sum_{z\in Z} \psi_{z,t,i,\omega} \qquad \forall s_d \in S_d, \forall t \in T, \forall i \in I, \forall \omega \in \Omega \qquad 4.46$$

$$\eta_{s_d,t,i,\omega} \geq \sum_{z\in Z}\left(\psi_{z,t,i,\omega} + \sum_{k\in K} \nu_{s_d,z,k,i,\omega}\right) - 1 \qquad \forall s_d \in S_d | s_d \geq 2\,, \forall t \in T, \forall i \in I, \forall \omega \in \Omega \qquad 4.47$$

$$\eta_{1,t,i,\omega} \geq \psi_{z,t,i,\omega} + \sum_{k \in K} \nu_{1,z,k,i,\omega} - 1 \qquad \forall t \in T, \forall z \in Z, \forall i \in I, \forall \omega \in \Omega \qquad 4.48$$

$$\phi_{s,t,i,\omega} \leq \eta_{s,t,i,\omega} \qquad \forall s_d \in S_d, \forall t \in T, \forall i \in I, \forall \omega \in \Omega \qquad 4.49$$

$$\phi_{s,t,i,\omega} \leq \xi_{t,i,\omega} \qquad \forall s_d \in S_d, \forall t \in T, \forall i \in I, \forall \omega \in \Omega \qquad 4.50$$

$$\phi_{s,t,i,\omega} \geq \eta_{s,t,i,\omega} + \xi_{t,i,\omega} - 1 \qquad \forall s_d \in S_d, \forall t \in T, \forall i \in I, \forall \omega \in \Omega \qquad 4.51$$

Finally, the clinical actions taken given the screening outcomes are enforced by Eqns. 4.52 to 4.56. The first of the expressions, 4.52, is a logical constraint requiring individual i to be at a single health state at every age. Equation 4.53 enforces the resulting lifetime, $Y_{s,t,i,\omega}$, values for the healthy state, $s = 1$. The only way for an individual to return to a healthy state after they progress through the various health states is if a screening action successfully detects an adenoma, $s = 2$. Constraint in Eqn. 4.54 removes an individual from an adenoma or undiagnosed CRC state, $s = 2, \dots 6$. This constraint modifies the resulting lifetime removing the further progression of the disease past state s if a screening test detects a non-healthy state at or before the given state. The clinically detected CRC health states are assigned in Eqn. 4.55. The states associated with clinically detected CRC, $s = 7, \dots, 10$, can only be reached if the associated undiagnosed CRC state is detected with a screening action, but these states can be avoided if a screening action detects an earlier health state. Lastly, Eqn. 4.56 enforces the deceased state for each individual.

$$\sum_{s \in S} Y_{s,t,i,\omega} = 1 \qquad \forall t \in T, \forall i \in I, \forall \omega \in \Omega \qquad .52$$

$$Y_{1,t,i,\omega} \leq \theta_{1,t,i}\left(1 - \xi_{t,i,\omega}\right) + \eta_{2,t,i,\omega} - \phi_{2,t,i,\omega} \qquad \forall t \in T, \forall i \in I, \forall \omega \in \Omega \qquad .53$$

$$Y_{s,t,i,\omega} \le \theta_{s,t,i}\left(1 - \xi_{t,i,\omega} + \sum_{s'=2}^{s}\left(\phi_{s',t,i,\omega} - \eta_{s',t,i,\omega}\right)\right) \quad \forall s \in [2,3,4,5,6], \forall t \in T, \forall i \in I, \forall \omega \quad .54$$

$$Y_{s,t,i,\omega} \le \theta_{s,t,i}\left(1 - \xi_{t,i,\omega} + \sum_{s'=2}^{s_d}\left(\phi_{s',t,i,\omega} - \eta_{s',t,i,\omega}\right) + \eta_{s_d,t,i,\omega} - \phi_{s_d,t,i,\omega}\right) \quad \forall (s, s_d) \in [(7,3),(8,4),(9,5),(10,6)], \forall t \in T, \forall i \in I, \forall \omega \in \Omega \quad .55$$

$$Y_{11,t,i,\omega} \ge \xi_{t,i,\omega} \quad \forall t \in T, \forall i \in I, \forall \omega \in \Omega \quad .56$$

To obtain the full models, referred to as CRCNL, CRCL1, and CRCL2, we must first add Eqn. 4.30 as a constraint to define the utility function for the model. Then, we replace the general form of the fixed recourse constraints, Eqn. 4.4, with our case study specific constraints, Eqn. 4.31 to 4.56. These changes yield three models, CRCNL, the nonlinear model, given in Eqns. 4.57 and 4.58, CRCL1, the linear model using the approach in Section 4.1.1.2.1, in Eqns.4.59 and 4.60, and CRCL2, linear model using the approach in Section 4.1.1.2.2, in Eqns.4.61 and 4.62.

$CRCNL$:

$$max \sum_{i \in I} \mathbb{P}_i \sum_{\omega \in \Omega}\left(\prod_{z \in Z}\sum_{j \in J}\sum_{s \in S}\sum_{t \in T} \mathbb{P}_{s,j}\left(\Xi_{z,\omega}\right)\theta_{s,t,i}X_{j,t,z}\right)F(\boldsymbol{X}, \boldsymbol{Y}_{i,\omega}, \boldsymbol{\theta_i}, \omega) \quad 4.57$$

$$s.t.\ \ 4.1, \mathbf{4.2}, 4.7, 4.8, 4.30 - 4.56 \quad 4.58$$

CRCL1:

$$max \sum_{i \in I} \mathbb{P}_i \sum_{\omega \in \Omega} \sum_{\substack{(s_1,j_1),\ldots,\\ (s_{|Z|},j_{|Z|}) \in (S \times J)^{|Z|}}} \mathbb{P}_{(s_1,j_1),\ldots,(s_{|Z|},j_{|Z|}),\omega} U^{L1}_{(s_1,j_1),\ldots,(s_{|Z|},j_{|Z|}),i,\omega} \quad 4.59$$

$$s.t. 4.1, \mathbf{4.2}, 4.7 - 4.15\ , 4.30 - 4.56 \quad 4.60$$

CRCL2:

$$max \sum_{i \in I} \mathbb{P}_i \sum_{\omega \in \Omega} U^{L2}_{|Z|,i,\omega} \quad 4.61$$

$$s.t. 4.5, 4.20, 4.21, 4.23 - 4.27\ , 4.30 - 4.56 \quad 4.62$$

4.1.1.4 Results and Discussion

To evaluate the three models, we solved each model by varying the total number of individuals in the cohort, $|I| = 100, 250, 500, 1000$, and actions within the timeframe, $|Z| = 1, 2, 3, 4$. When varying the population, the number of screening actions was fixed to two, and while varying the number of screening actions, the size of the population was set to 100. The values of the uncertain parameters, $\boldsymbol{\theta}$, $\boldsymbol{\tau}^{Nat}$, and $\boldsymbol{\tau}^{CRC}$**,** were generated sampling from the constructed empirical distributions. The model instances were constructed in Python 3.8 using Pyomo 6.1.2. Each model was solved directly, CRCNL was solved using BARON 21.1.13, and CRCL1 and CRCL2 were solved using CPLEX 20.10. The model instances were solved on an HPC cluster node with two Intel® Xeon 6248R 3.0 GHz processors with 48 cores and an available 384 GB of RAM. The maximum wall time was 24 hours.

4.1.1.4.1 Model Performance

A summary of the model size and solution wall time for CRCNL are in Table 4.2, while the statistics for CRCL1 and CRCL2 are in Table 4.3 and Table 4.4, respectively. In the CRCSP, binary variables define the first and second stage decisions, and the only continuous variables in the model track the utility for each individual-screening outcome pair. As the base model, CRCNL shows the general growth of the model size. As seen in Table 4.2, the model size grows linearly as the cohort size, $|I|$, increases while holding the total number of screens in a lifetime, $|Z|$,

constant. When the cohort size is held constant and the number of screens in a lifetime are varied, model size grows exponentially. The global solver BARON did not yield the optimum solution within the 24-hour wall time limit for any instance of CRCNL, revealing the intractability of the model, which stems from its size, nonlinearities in its objective function, and combinatorial complexity due to the shear number of binary variables.

Table 4.2. Model statistics for CRCNL across all instances.

Instance ($\|\boldsymbol{Z}\|, \|\boldsymbol{I}\|$)	Total Variables	Binary Variables	Continuous Variables	Constraints	Solution Wall Time (h:m:s)
(2,100)	866,962	866,562	400	2,422,483	24:00:00
(2,250)	2,167,162	2,166,162	1,000	6,056,083	24:00:00
(2,500)	4,334,162	4,332,162	2,000	12,292,083	24:00:00
(2,1000)	8,668,162	8,664,162	4,000	24,584,083	24:00:00
(1,100)	411,281	411,081	200	1,140,201	24:00:00
(2,100)	866,962	866,562	400	2,422,483	24:00:00
(3,100)	1,822,643	1,821,843	800	5,128,965	24:00:00
(4,100)	3,822,724	3,821,124	1,600	10,825,847	24:00:00

The linearization approach used to construct CRCL1, presented in Section 4.1.1.2.1, from CRCNL adds both continuous and binary variables to the model. The model size, which grew exponentially with the total number of screens in a lifetime, $|Z|$ for CRCNL, grows even more rapidly for CRCL1. The linearization of CRCL1 increases the model variables from 1.01 to 14.02 times the number of variables compared to CRCNL and 1.01 to 10.74 times the number of constraints when compared to the original CRCNL, depending on the instance. The instance $(|Z|, |I|) = (4,100)$ of CRCL1 was not solved to optimality due to its space complexity, resulting from its model size, i.e., resulted in an out-of-memory error. Enumerating all potential outcomes of the screening actions in CRCL1 made the model computationally intractable, even at a low

number of individuals within the cohort, $|I|$. The rapid growth in model size is one drawback when using this method for linearizing the model. Despite the rapid increase in model size, CPLEX solved the remaining instances of CRCL1 within the allowed wall time. The solution time of CRCL1 showed an exponential increase when increasing the cohort size and the number of screens in a lifetime. Much like the model size, the solution time increases more quickly as the number of screens in a lifetime increases rather than the increase in cohort size

Table 4.3. Model statistics for CRCL1 across all instances.

Instance $(\|\boldsymbol{Z}\|, \|\boldsymbol{I}\|)$	Total Variables	Binary Variables	Continuous Variables	Constraints	Solution Wall Time (h:m:s)
(2,100)	927,462	878,662	48,800	2,664,483	00:03:15
(2,250)	2,318,412	2,196,412	122,000	6,661,083	00:15:01
(2,500)	4,636,662	4,392,662	244,000	13,322,083	01:11:56
(2,1000)	9,273,162	8,785,162	488,000	26,644,083	08:06:20
(1,100)	414,581	412,181	2,400	1,151,201	00:00:44
(2,100)	927,462	878,662	48,800	2,664,483	00:03:15
(3,100)	3,020,543	1,954,943	1,065,600	10,186,765	05:53:22
(4,100)	53,602,124	5,285,224	23,427,200	16,241,047	OOM

The linearization approach used to construct CRCL2 introduces only continuous variables and produces a slightly larger model than CRCNL. The size of CRCL2 increases about 1% in both the number of variables and constraints compared to CRCNL when increasing the value of $|I|$, a behavior shared by CRCL1. When the value of $|Z|$ is increased, the size of CRCL2 only grows between 0.6 to 2% in the number of variables and 0.8 to 2.7% in the number of constraints compared to CRCNL. All instances of CRCL2, but one, were solved to optimality by CPLEX. The instance that could not be solved is $(|Z|, |I|) = (4,100)$. However, CPLEX was not able to solve

this instance of CRCL2 due to reaching the maximum allowed wall time, unlike the CRCL1 case which had a space complexity.

Table 4.4. Model statistics for CRCL2 across all instances.

Instance $(\|\boldsymbol{Z}\|, \|\boldsymbol{I}\|)$	Total Variables	Binary Variables	Continuous Variables	Constraints	Solution Wall Time (h:m:s)
(2,100)	876,562	866,562	10,000	2,458,483	00:04:02
(2,250)	2,191,162	2,166,162	25,000	6,146,083	00:21:30
(2,500)	4,382,162	4,332,162	50,000	12,292,083	02:27:39
(2,1000)	8,764,162	8,664,162	100,000	24,584,083	19:28:31
(1,100)	413,681	411,081	2,600	1,149,201	00:00:45
(2,100)	876,562	866,562	10,000	2,458,483	00:04:02
(3,100)	1,851,443	1,821,843	29,600	5,236,965	17:46:48
(4,100)	3,899,524	3,821,124	78,400	11,113,847	24:00:00

As expected, the linear models, CRCL1 and CRCL2, have lower space and time complexity than the nonlinear model, CRCNL. However, each model has strengths and drawbacks stemming from the linearization technique employed. Figure 4.1 presents a plot comparing the solution time versus the model size for CRCL1 and CRCL2. The dotted line in Figure 4.1 shows the maximum wall time; any points sitting higher than that line are the instances that were not solved to optimality, whether due to space or time complexity. The data point labels are the model instances, $(|Z|, |I|)$. For both models, the solution time grows steadily as the model size increases due to increasing $|I|$, with the growth for CRCL2 being slightly higher than that of CRCL1. On the other hand, the model size of CRCL1 is slightly larger than CRCL2. The solution time grows quicker with increasing $|Z|$ for both CRCL1 and CRCL2. The growth is faster for CRCL2, which leads the solution time to exceed the limit for instance $(|Z|, |I|) = (4,100)$. The model size shows a similar growth trend, increases in $|Z|$ yields a larger model size than increases in $|I|$, with CRCL1

model size growing much faster than CRCL2. The CRCL1 model scales poorly for increases in $|Z|$ but can be solved more quickly. While the solution times for CRCL2 model grow quicker than CRCL1, as noted by Medal et al. (2016), this formulation keeps the model compact with increasing $|Z|$ when compared to CRCL1.

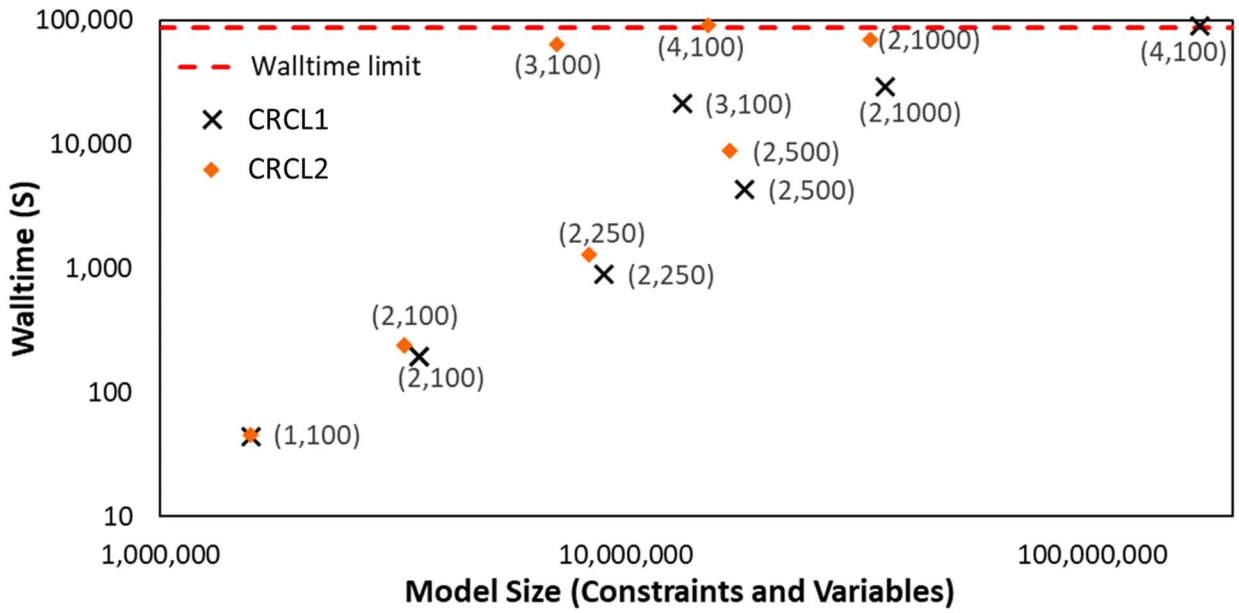

Figure 4.1 Solution time versus model size for CRCL1 and CRCL2.

Various metrics were investigated to understand the impact of each linearization procedure on the problem. These metrics, outlined in Figure 4.2, are the initial gap of the problem (Figure 4.2-A), solution wall time of the root relaxation problem (Figure 4.2-B), the total number of nodes created during the branching step of CPLEX (Figure 4.2-C), and the node at which the optimal solution was found (Figure 4.2-D). Each plot compares one metric to the solution time for CRCL1 and CRCL2 instances, identified by the labels. Figure 4.2-A reveals that CRCL1 has a smaller initial percent gap than CRCL2 for nearly every instance. However, for instance (3,100), the initial gap for CRCL2 is nearly half that of CRCL1, while the solution time for CRCL2 is nearly three times longer than CRCL1. Similar trends are observed for the total number of nodes explored and

the node number in which CPLEX identifies the optimal solution (Figure 4.2 C and D). All instances except (2,100) and (3, 100) show that CPLEX generated fewer branches for CRCL1 than CRCL2 before proving optimality. These trends partially explain the discrepancy in the solution time, but instance (3, 100) shows that these metrics alone do not fully explain the difference in solution times. Additionally, the root relaxation solution time plot, Figure 4.2 B, shows that the linear programming (LP) relaxations of CRCL2 is solved faster than the LP relaxations of CRCL1 for more than half of the cases.

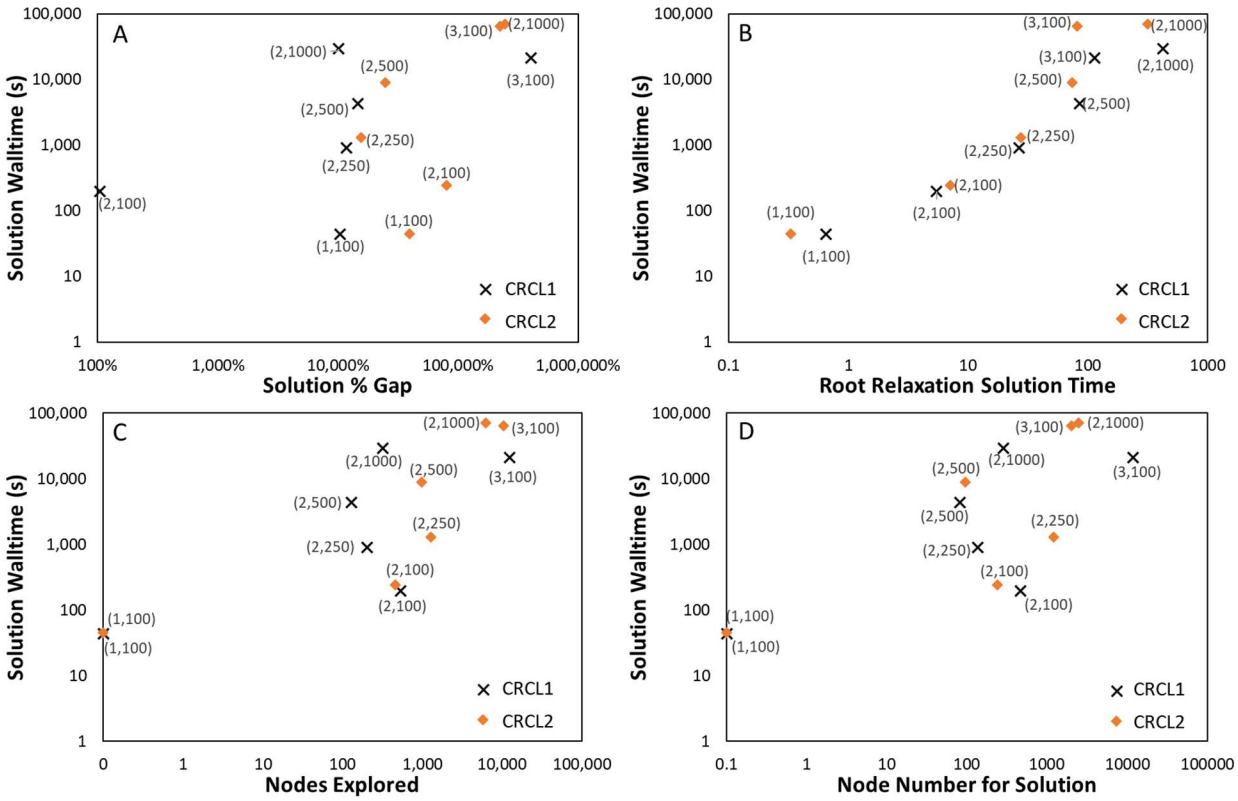

Figure 4.2 Model statistics comparison of CRCL1 & CRCL2 showing A) initial solution gap versus overall solution time, B) root relaxation time versus overall solution time, C) nodes explored versus overall solution time D) node number for identifying optimal solution versus overall solution time.

The LP relaxations of both models were further studied to understand the differences in solution time. The metrics considered for this analysis are given in Table 4.5, and they include the

solution time, iterations needed to solve the LP, and the optimum objective function value. The LP relaxations of CRCL2 are solved much quicker than CRCL1, and CPLEX solves them in fewer iterations than CRCL1. The solution times for the LP relaxations of CRCL2 ranged from about half to a quarter of those for the LP relaxation of CRCL1. However, based on the optimum objective function values of the relaxed models, CRCL1 is a much tighter formulation than CRCL2, where the optimal value of relaxed CRCL1 is half that of CRCL2, yielding faster solution times for CRCL1.

Table 4.5 Solution statistics for the LP relaxation of CRCL1 and CRCL2.

Case	Solution Walltime (s)		Iterations		Objective Value	
	CRCL1	CRCL2	CRCL1	CRCL2	CRCL1	CRCL2
(2, 100)	11.75	5.76	39,985	4,769	166.00	320.00
(2, 250)	135.32	23.45	130,092	17,149	162.99	314.88
(2, 500)	335.28	100.67	337,447	36,433	162.86	316.16
(2, 1000)	1141.21	291.57	603,255	74,645	161.59	314.88
(1, 100)	2.44	1.97	1,267	1,450	80.00	159.99
(3,100)	78.42	30.66	156,019	31,829	386.75	640.00

4.1.1.4.2 Solution

The optimal solutions and the expected QALY gained per person for the different instances of the CRCSP are presented in Figure 4.3. Figure 4.3-A displays the change in the expected QALY gained as a function of the total screenings tests in a lifetime, and Figure 4.3-B shows the change in the expected QALY gained as a function of the change in total cohort size. The data labels depict the age(s) at which the eligible members of the cohort receive a colonoscopy. From Fig. 4.3-A, it can be seen that there is a diminishing return for adding more screens in a lifetime. From one to two screens in a lifetime, there is a 9% increase in expected QALY gained, while from two to three screens, the increase is only 0.1%.

The ages for screening (Figure 4.3-A) differ significantly from the current recommendations, screening with a colonoscopy every ten years starting at age 50. However, the additions of new screening tests do show an expected trend of the timing of the screening ages, though the ages of the actual screens were skewed towards later in life. For one screening in a lifetime ($|Z| = 1$), it was expected to be recommended near the middle of the age range to catch a bulk of the early CRC cases and the larger adenomas to achieve a high expected gain in QALY. With more screening, the timing of the first screen can shift to an earlier age to capture the cases where an adenoma forms earlier in life that are less likely than a majority of the CRC cases, but provide a very high potential for gain in QALY. While the final screen can be later in life to capture the cases where CRC is developed but shows little gain in QALY.

The discrepancy of the screening ages seen in Figure 4.3-A, when compared to the current recommendations, is shown to improve, Figure 4.3-B, at the introduction of larger cohort sizes, $|I|$. As the cohort size increased there is a shift for the screening ages to be earlier in life and, with that, a higher value for the expected gain in QALY for the cohort. The occurrence of CRC in the general population, though one of the most common forms of cancer, is still relatively rare, where only around 4-5% of men develop CRC within their lifetime[178]. The prevalences of CRC in the cohorts of the case study are 2%, 5.2%, 4.6%, and 4.7% for the values of $|I|$ at 100, 25, 500, and 1000, respectively, which explains the increase in the expected QALY gained for the optimal solutions, as there is an increase in the population that has CRC. In addition to the prevalence, the disease progression impacts the expected gain in QALY. The average ages for cancer development in the different instances are 82.5, 72.3, 69.8, and 70.6 for the values of $|I|$ at 100, 25, 500, and 1000, respectively. The larger number of individuals that develop CRC in their lifetime allows for

a wider range of potential CRC progression timelines, better representing the uncertainty space where screening benefits an individual.

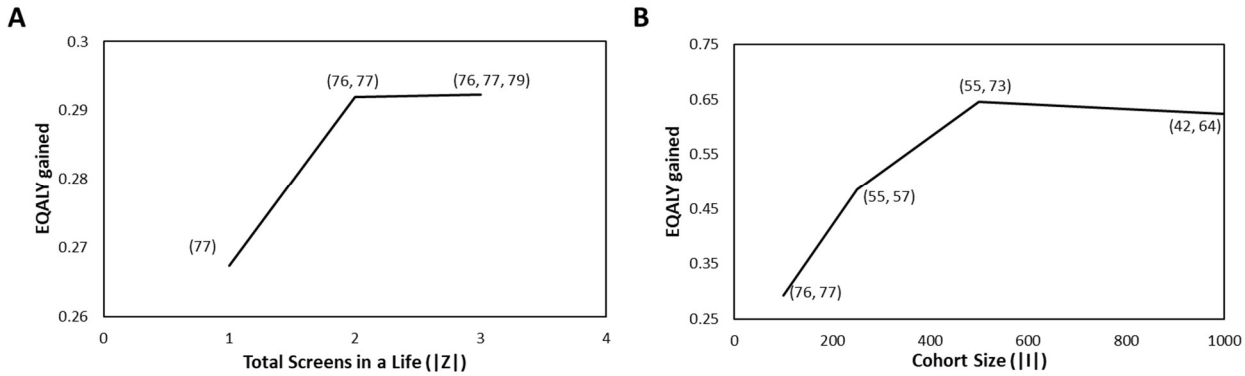

Figure 4.3 EQALY gained with A) total screens in a lifetime and B) the size of the cohort

4.1.2 Efficient Scenario Generation for TSSP

A major drawback of the case study presented in Section 4.1.1 is the large changes in both optimal solution and objective value as the cohort size, $|I|$, increases. With CRC being a relatively rare event, traditional naïve Monte Carlo methods for generating SP scenarios were not able to fully represent the uncertainty space without requiring a large sample size and, in turn, leading to larger model sizes. The following study aims to assess which widely applicable scenario construction methods can efficiently generate samples that best represent the rare event space for SPs where that rare-event space primarily drives the optimal decisions, and the uncertain distribution is either continuous or discrete and infinite.

4.1.2.1 Scenario construction methodologies

4.1.2.1.1 Monte Carlo Methods

The Monte Carlo method (MCM) is an umbrella term encompassing many different algorithms that all rely heavily on generating random samples from a distribution. The most

simplistic of these methodologies is known as the crude, or naïve, Monte Carlo (CMC) method. In naïve Monte Carlo, samples are generated directly from the distribution under study, assuming that given enough samples, the distribution can be described by a discrete set of realizations. In addition to naïve Monte Carlo, other Monte Carlo based methodologies have been created with uses ranging from variance reduction methods[179] to distribution estimations[180]. For this study, we focus on using two different MCMs, crude Monte Carlo and importance sampling (IS). For clarity, when referring to different MCMs (CMC versus IS), we will refer to the distribution used to sample from, e.g. naïve Monte Carlo is referred to as sampling from the original distribution. When referring to the sampling approach of generating samples directly from a specified distribution, we will refer to that as CMC sampling.

Importance sampling (IS) is a specialized MCM used for variance reduction. The samples in IS are not directly generated from the original distribution under study. Instead, they are sampled from an auxiliary distribution, known as the importance distribution (ID). An importance distribution is used to generate samples from a region of interest, usually an area with low probability but a high impact on the system under study. The region of interest in theID has a higher probability. The samples are then reweighted by a likelihood ratio to maintain the original distribution properties. The likelihood ratio, $w(x)$, is presented in Eqn. 4.63, where $p(x)$ is the original probability density function (PDF) of the distribution under study, and $ID(x)$ is the PDF of the importance distribution.

$$w(x) = \frac{p(\boldsymbol{x})}{ID(\boldsymbol{x})} \tag{4.63}$$

Though any distribution can be used for the ID, the selection or construction of the ID is not a trivial task. If a poorly suited ID is used, the intended variance reduction behavior of IS can actually cause an increase in variance, providing no benefit over random sampling[181]. The optimal ID for variance reduction is presented in Eqn. 4.64[181], where $f(\boldsymbol{x})$ is referred to as the importance function where $f: \boldsymbol{x} \to \mathbb{R} | f(\boldsymbol{x}) \geq \boldsymbol{0}\ \forall \boldsymbol{x} \in \boldsymbol{\mathcal{X}}$, and $\mathbb{E}_p$ is the expectation operator over the distribution $p(\boldsymbol{x})$.

$$ID(x) = \frac{p(x)f(x)}{\mathbb{E}_p[f(x)]} \tag{4.64}$$

For this work, we construct the ID using the methodology of Papavasiliou and Oren (2013). We estimate the value of $\mathbb{E}_p[f(\boldsymbol{x})]$ by generating a large number, M, of CMC samples from our original distribution. We then generate N sub-samples from the large set of M samples through sampling from the estimated ID.

4.1.2.1.2 Quasi-Monte Carlo and Space-filling Methods

As with importance sampling, Quasi-Monte Carlo (QMC) and space-filling sampling methods were conceived as ways to improve the generation of samples for computational experiments. Quasi-Monte Carlo methods generate samples using a low-discrepancy sequence. These sequences mimic independent random samples used in CMC sampling but have been constructed to ensure quick and even coverage of the sampling domain. They are deterministic and have been shown to provide more accurate results in numerical integration and moment estimation than CMC sampling[183]. Of the number of low-discrepancy sequences, our study employs both Sobol' and Halton sequences.

As its name suggests, space-filling methods generate samples to minimize gaps within the sampling region. Example space-filling approaches include Latin Hypercube Sampling (LHS), minimax/maximin distance sampling, and maximum projection sampling[184]. Our study uses LHS. Latin hypercube sampling divides each distribution of a d-dimensional hypercube into N bins of equal probability. It then generates N samples ensuring no two samples are in the same bin on any axis of the hypercube[96].

4.1.2.1.3 Clustering Methods

Clustering methodologies are a group of unsupervised classification algorithms used to group or cluster data that share similar features. There are different approaches to defining similarities between data points, including hierarchical clustering, centroid-based clustering, or density-based clustering. Unlike the previously presented methods, clustering is a data reduction approach. To use this methodology for scenario generation, a representative point from each cluster is used as a single sample point or scenario. We focus on the use of centroid-based clustering. Centroid-based clustering defines each cluster through a representative center point and assigns points of data to a cluster based on the nearest cluster center.

The two centroid-based clustering algorithms used in this study are the k-means and x-means algorithms. The k-means algorithm determines the location of the specified number of centroids by iteratively assigning data points to a cluster given the current location of the centroids and then updating the centroids by calculating the average location of the data points in the cluster. This process is repeated until the centroids converge to stable points[185]. A drawback of the k-means algorithm is its requirement of the number of clusters as an input. The correct number of clusters for the data set may not be known *a priori*. To determine the correct number of clusters, we employed the elbow method with the sum of the squared distance (SSD) from the nearest cluster

center as the metric to determine the number of clusters, k[186]. This method tends to be subjective as it is up to the user to decide when the trend begins to show diminishing returns on the addition of new cluster centers. To avoid this subjectivity, we use the point on the SSD curve with the maximum distance to a line connecting the minimum and the maximum number of clusters[187].

The x-means algorithm is an extension of the traditional k-means algorithm, where the algorithm determines the number of clusters. The algorithm introduces new clusters by subdividing existing clusters and evaluating if the subdivisions improve the clustering results[188]. The Bayesian Information Criterion (BIC) criterion was used to evaluate the improvements. The BIC was calculated before and post subdividing the clusters. The clustering model with the greater BIC is selected as the best.

4.1.2.2 Case Study: Colorectal Cancer Screening Problem (CRCSP)

4.1.2.2.1 Problem Statement and Description

The colorectal cancer (CRC) screening problem (CRCSP) aims to identify the optimal screening decisions to maximize the expected quality-adjusted life-years (QALY) gained for a cohort of individuals. This problem is formulated as a two-stage SP (TSSP) with fixed recourse actions, with the first-stage decisions being the screening decisions for the entire cohort and the recourse actions being the clinical actions taken given the outcomes of the screening tests. The screening decisions determine the ages and the screening tests to use for a cohort-wide screening strategy, for a predefined number of screening tests taken within a lifetime. The objective function calculates the gain in QALY by comparing the difference between a lifetime where screening occurs and a lifetime where screening does not occur for each individual. The full formulation of this model can be found in Eqns. 4.59 and 4.60.

In this problem, there are two sources of uncertainty: 1) the progression of CRC within an individual to describe their lifetime ($\boldsymbol{\theta}$) and 2) the outcomes of the imperfect screening tests (Ξ). However, the second source of uncertainty can be fully elucidated due to the binary nature of the testing outcomes. A test either detects CRC or its precursors (adenomas) or does not. This property of the second source leaves the first source of uncertainty as the focus of the study. The uncertainty distribution for $\boldsymbol{\theta}$ is constructed using data from a microsimulation model that simulates the progression of CRC[157,170]. Each point of this data represents a unique realization of the progression of cancer, or lack thereof, within an individual and is used to describe the progression through a number of health states. A total of six features are used to capture the simulated lifetime of CRC within an individual for a single data point; these are 1) the age at which an adenoma (precursor to CRC) first develops, 2) the time between the first developed adenoma and the first case of cancer development, 3) the time for cancer to be detected by symptomatic means, 4) the stage of the symptomatically detected cancer, 5) the number of years survived after cancer detection, and 6) the age at which the individual would die from sources outside of CRC.

4.1.2.2.1.1 Methodology

To assess the scenario generation methodologies, instances of the toy problem were constructing using 128, 512, and 1,024 realizations for uncertain parameters $\boldsymbol{\tau^F}$ and $\boldsymbol{\tau^W}$ for each of the sampling methodologies, CMC sampling, LHS, Sobol, and Halton. Due to the random nature of CMC sampling and LHS, thirty iterations of each scenario instance was used to obtain an average performance of these methodologies. For both Sobol and Halton sequences the first 1,024 samples of the sequences were skipped. In addition to the scenario generation methodology, the two Monte Carlo methodologies used to generate the scenarios, sampling directly from the original uncertain distribution and sampling from an ID, were also evaluated. To better evaluate the IS

approach, the ID was constructed with 100,000, 1,000,000, and 2,000,000 CMC samples from the original distribution, with the scenarios generated from each ID using each sampling methodology described previously.

The importance function (IF), $f(\boldsymbol{\theta})$, used to construct the ID is defined as the ratio of the quality life years an individual would have lived in the absence of CRC and the quality life years lived based on the progression of CRC. Using this definition yields $f(\boldsymbol{\theta}) = 1$ if an individual never develops CRC within their lifetime and $f(\boldsymbol{\theta}) > 1$ if an individual develops CRC. For the scenario reduction methods, k-means and x-means algorithms, data sets with 100,000, 1,000,000, and 2,000,000 points randomly drawn, without replacement, from the data used to construct the uncertainty distribution for $\boldsymbol{\theta}$ are generated. The data sets are then provided to the clustering algorithms with the six features outlined in the previous section and seven features, with the additional feature being the importance function value of each point when ID is used. The total number of screening tests within a lifetime was set to two, with only one screening test option available for the screening decisions. The CRCSP model was implemented within Python 3.8 using Pyomo 6.1.2. All model instances were solved using CPLEX 20.10 on a single HPC node using two 3.00 GHz Intel® Xeon Gold 6248R processors for up to 48 cores with 384 GB of RAM.

4.1.2.2.2 Results

A simulation-optimization approach using the CRC microsimulation was used to generate the optimum solutions for assessing the quality of the TSSP solutions obtained with the scenarios constructed using the various methods. The microsimulation model was the one that was used to generate the data for the uncertain parameter distributions. We modified the microsimulation model to follow the assumptions of the TSSP model. A simulation optimization determined the optimal ages for screening using the global derivative-free optimization (DFO) solver

TOMLAB/glcSolve[90] and the local search solver Derivative Free Line-search (DFL)[81]. The maximum number of simulation evaluations was set to 5,000. Upon termination, the DFL[81] was used to confirm the local optimality of the solution generated by TOMLAB/glcSolve. Figure 4.4 presents the solutions yielded by the two DFO solvers. Each data point in Figure 4.4 depicts the ages for the two screening tests for the population and the expected gain in QALY for implementing that screening strategy. The best-identified solution, $\boldsymbol{x}_{DFO}^{*}$, and the resulting objective function value, $\boldsymbol{y}_{DFO}^{*}$, is labelled and represented by a star, and the remaining points of data go from a lighter to darker shade as the expected QALY increases.

We used two metrics to assess the quality of the TSSP solutions. The first metric is the Euclidian distance, d_{opt}, between the TSSP solution, $\boldsymbol{x}_{SP}^{*}$, and the best solution identified by the DFO, as shown in Eqn. 4.65.

$$d_{opt}(\boldsymbol{x}_{DFO}^{*}, \boldsymbol{x}_{SP}^{*}) = \sqrt{(\boldsymbol{x}_{DFO}^{*} - \boldsymbol{x}_{SP}^{*})^2} \tag{4.65}$$

The second metric is the percent difference between the objective function values of the TSSP solution and the best solution identified by the DFO, $\%\ err$. The calculation is presented in Eqn. 4.66, where y_{SP}^{*} is the optimal objective function value of the TSSP.

$$\%\ err = \frac{y_{SP}^{*} - y_{DFO}^{*}}{y_{DFO}^{*}} \tag{4.66}$$

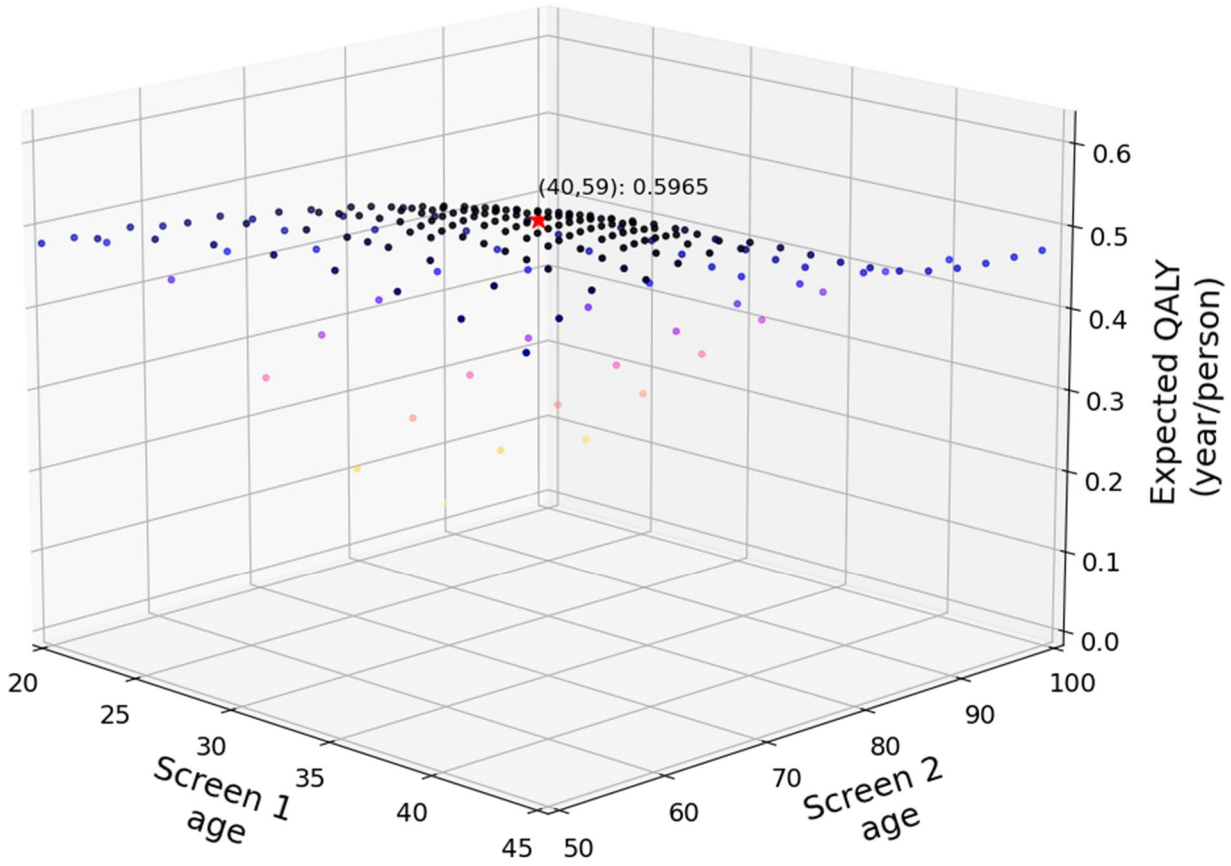

Figure 4.4 Simulation values of expected gain in QALY for two colonoscopies in a lifetime.

4.1.2.2.2.1 Generation of samples from the original distribution

The percent error and the distance from optimum metrics are displayed in Figure 4.5 for the sampling methods when using the original distribution. Figure 4.5-A shows the percent difference in the objective function values and Figure 4.5-B distance from the optimum. The whiskers in Figure 4.5 are the 5^{th} and 95^{th} percentiles of the metrics, and the diamond are the mean values for each sampling method. According to Figure 4.5-A, for CMC sampling and LHS, the average percent difference decreases as the number of scenarios increases. For Halton sampling, the percent difference increases with the scenario size, suggesting that with the addition of more scenarios, the benefits of screening are over-estimated. A similar trend is observed with Sobol

sampling, but the initial scenario set of 128 under-represents the benefits of screening, yielding a negative percent difference. Including more scenarios increases the percent difference. This behavior suggests that CMC sampling and LHS generate scenarios that do not bias the rare-event space compared to Sobol or Halton sequences for the CRCSP. According to Figure 4.5-B, the distance from the DFO solution decreases for all four sampling methods as the number of scenarios increases. This is due to an increase in the number of scenarios where CRC develops in an individual's life, more accurately representing the disease prevalence in the population. For scenarios generated using the Sobol and Halton sequences, the CRC cases are overrepresented in the scenario sets and may lead to an overestimate of screening benefit for those scenario sets. However, as the scenario size increases, the TSSP solutions get closer to the DFO solution for scenarios generated using the two sequences. The TSSP solutions obtained using the scenario sets constructed with LHS have a lower variance than those obtained using the scenarios constructed by CMC sampling for both percent error and distance metrics. Using the 1024 scenarios generated by the CMC sampling yields solutions closer to the DFO optimum than the ones generated by the LHS. However, the TSSP using the scenarios generated by the latter method approximates the optimal objective function value better than CMC sampling.

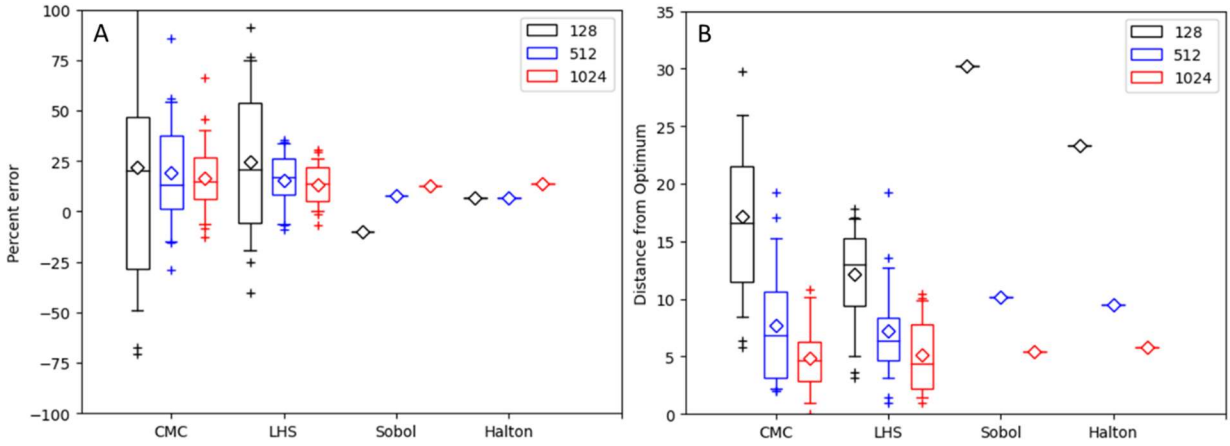

Figure 4.5 Comparison of the increase in number of scenarios for CMC sampling, LHS, and Sobol' and Halton sequence on the impact of the A) percent error and B) distance from the optimum when sampling from the original distribution.

All CRCSP TSSP solutions obtained using the scenario sets constructed by all approaches are presented in Figure 4.6. The solution density, defined as the percentage of solutions within a given distance from the optimal solution identified by the DFO, is plotted in Figure 4.7. The markers and lines are used to represent the different sampling methodologies, while the colors represent the total number of scenarios used in the scenario sets. The optimal solution is marked in Figure 4.6 using a triangle. As seen from Figure 4.5-B, there was a decrease in the average distance with an increase in the scenario size. This observation is better illustrated in Figure 4.6, where the spread of the TSSP solutions decreases, and the solutions shift closer to the optimal solution with increasing scenario size. Figure 4.7 further emphasizes this observation, where the increase in solution density shifts to lower distances, and the slope of the increase is greater for larger scenario sizes.

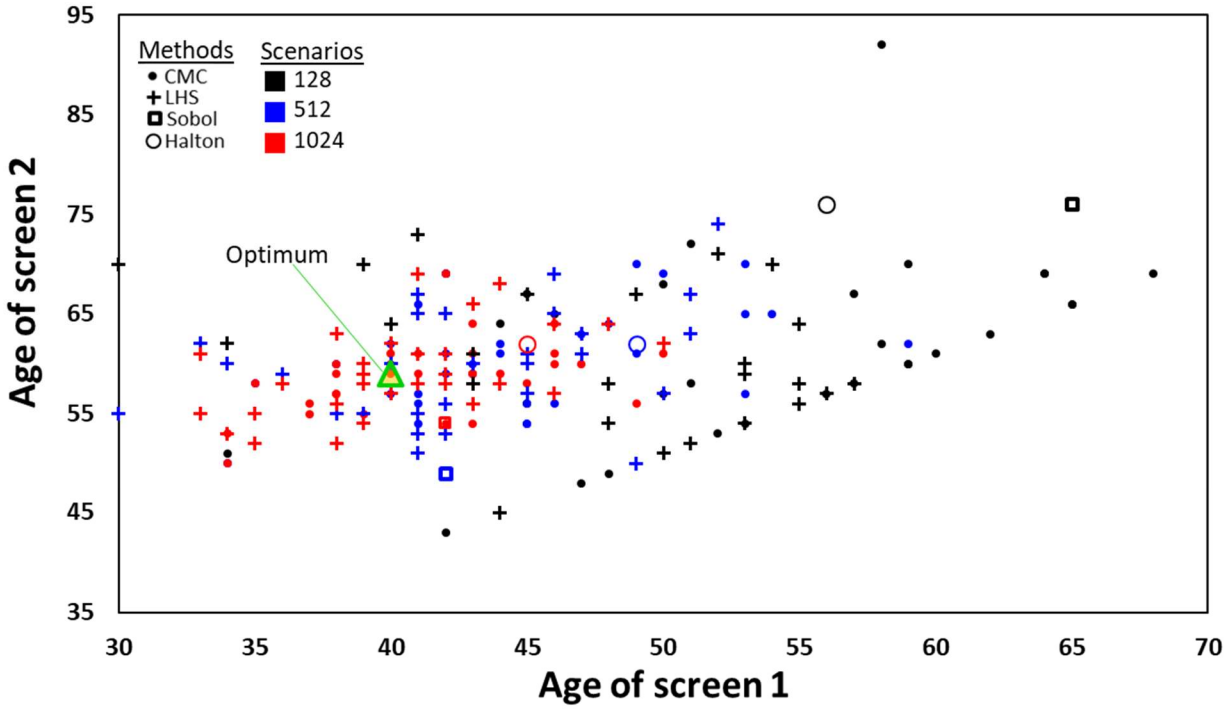

Figure 4.6 Optimal solutions generated from each sampling methodology for the original distribution for scenario set sizes of 128, 512, and 1024.

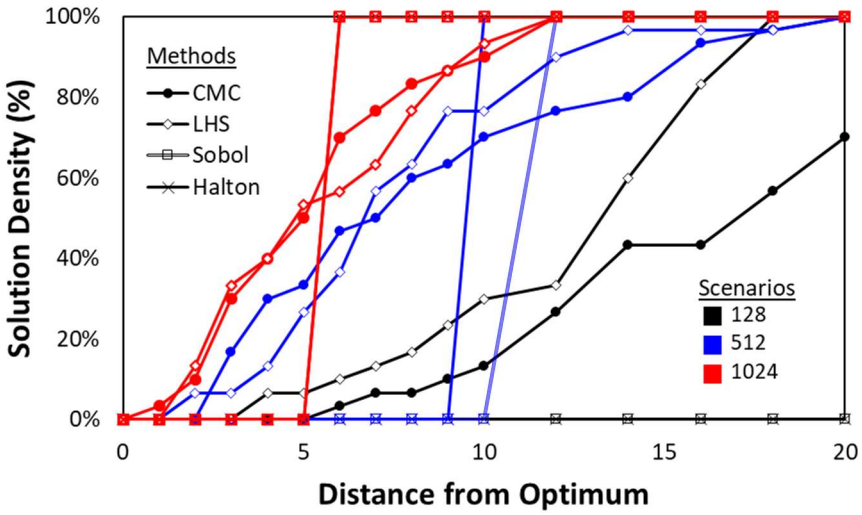

Figure 4.7 Solution density plot for sampling from the original distribution using different sampling methods for scenario set sizes of 128, 512, and 1024.

4.1.2.2.2.2 Generating samples from an importance distribution

The percent differences for TSSP solutions using scenarios generated by sampling from importance distributions are plotted in Figure 4.8 A-D. Each plot in Figure 4.8 corresponds to a

different sampling method, Figure 4.8-A CMC sampling, Figure 4.8-B the Halton sequence, Figure 4.8-C LHS, and Figure 4.8-D the Sobol sequence, used to generate the scenarios. The x-axis of the plots gives the number of samples used to construct the importance distributions (IDs). The colors correspond to the size of the scenario set. The trends observed with increasing scenario size in the previous section hold, for the most part. However, when Sobol and Halton sequences were used to generate scenarios from the ID constructed using 2,000,000 samples, an increase in scenario size led to a slight decrease in the percent difference from the DFO optimum objective function value rather than an increase.

No clear trends were observed between the percent difference variance for TSSP solutions obtained using scenarios generated with CMC sampling and LHS and the number of samples used to construct the IDs from which the methods generated the scenarios (Figure 4.8). This observation concurs with the expectation that the IS methodology reduced the variance of the samples. It is noted for CMC sampling, that IS has a more pronounced variance reduction effect on the solutions for 128 and 512 scenarios, whereas the impact of using an ID to generate for 1,024 scenarios is not as noticeable as seen by the spread of solutions in Figure 4.8-A. When LHS was employed with IS, a noticeable variance reduction was observed across all scenario sizes. For CMC sampling (Figure 4.8-A) and LHS (Figure 4.8-C), the average value of the percent difference was lower for scenarios generated using IS than the original distribution. The average percent difference however, is not statistically different when comparing the sampling distributions due to the variability of the solutions. Using the ID for generating scenarios with the Halton (Figure 4.8-B) and Sobol (Figure 4.8-D) sequences increased the percent difference compared to using the original distribution. We hypothesize that this increase results from over representation of the CRC cases in the scenarios, which was also observed with the original distribution.

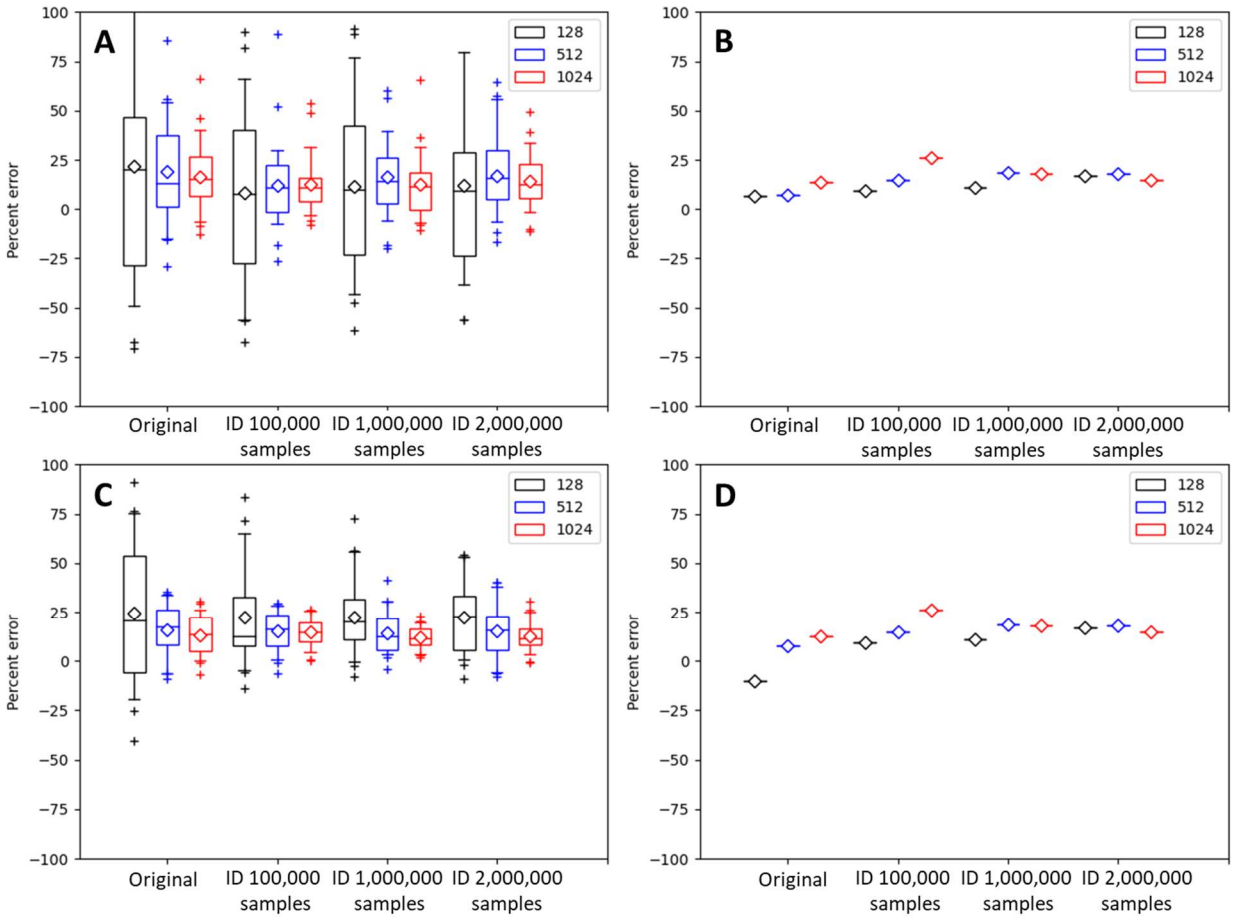

Figure 4.8 Percent errors of the sampling methods A) CMC sampling, B) Halton sequence, C) LHS, D) Sobol sequence comparing the distribution used to sample from.

FiguresFigure 4.9 - Figure 4.12 depict the solution density of the different sampling methodologies to compare the impact of the sampling distribution on scenario construction. For each figure, plot A shows the solution density for the TSSP model when using 128, 512, and 1,024 scenarios for all four sampling distributions, plot B for the model with 128 scenarios, plot C 512 scenarios, and plot D 1,024 scenarios. For all sampling methods, it is observed that generating the scenarios using an ID, regardless of the number of samples used to create the ID, generally provides a spread of TSSP solutions closer to the optimal than using the original distribution. This benefit is more prominently seen in the smaller scenario sizes. In 1,024 scenarios, there is little difference in the TSSP solutions between the different distributions.

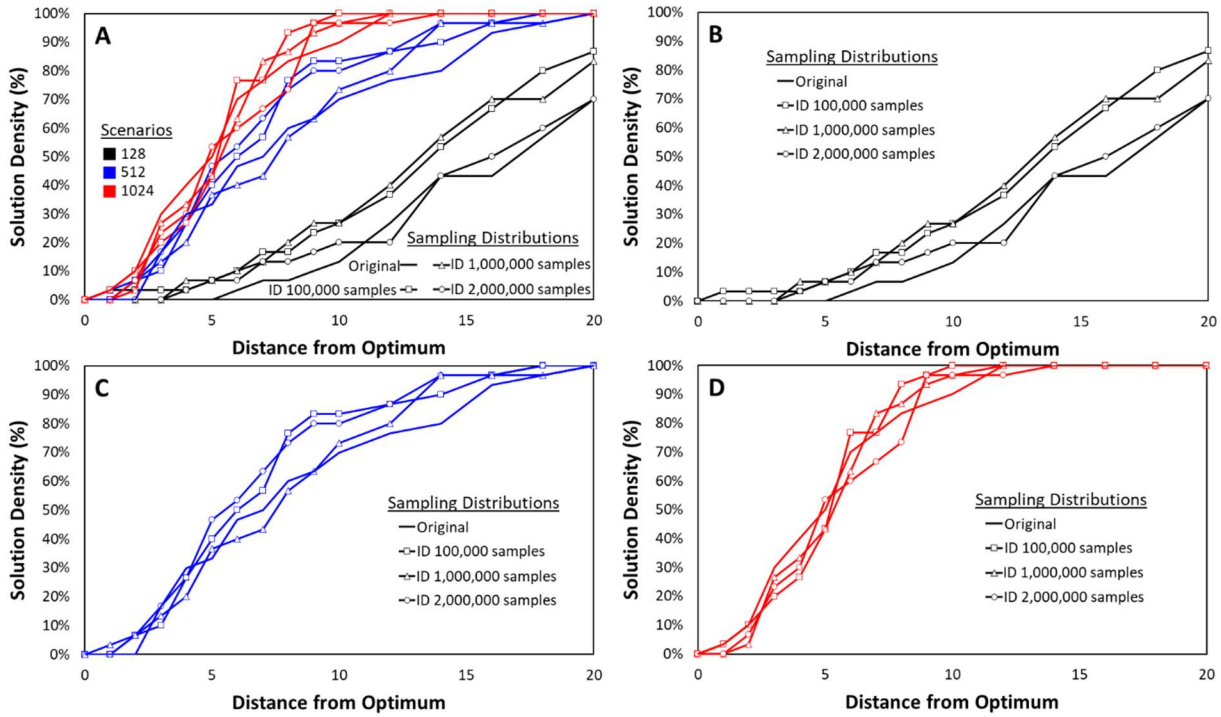

Figure 4.9 Solution density plot for CMC sampling comparing the distributions sampled from and the scenarios generated where A) displays 128, 512, and 1,024 scenarios, B) only displays 128 scenarios, C) only displays 512 scenarios, and D) only displays 1,024 scenarios.

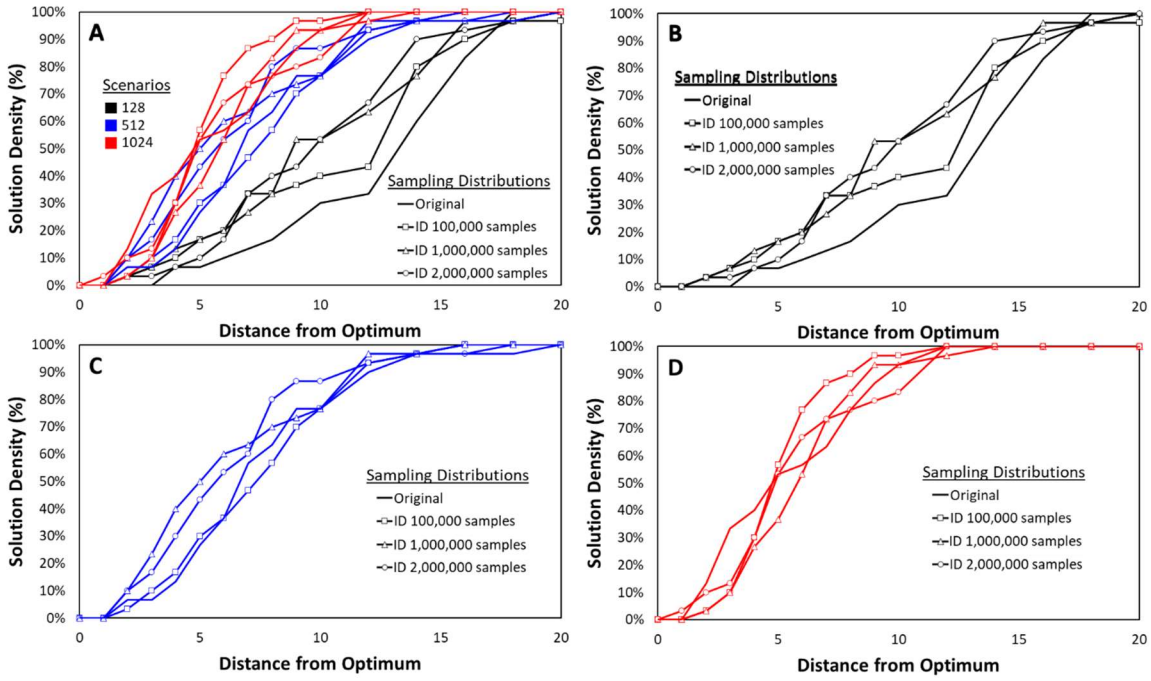

Figure 4.10 Solution density plot for LHS comparing the distributions sampled from and the scenarios generated where A) displays 128, 512, and 1,024 scenarios, B) only displays 128 scenarios, C) only displays 512 scenarios, and D) only displays 1,024 scenarios.

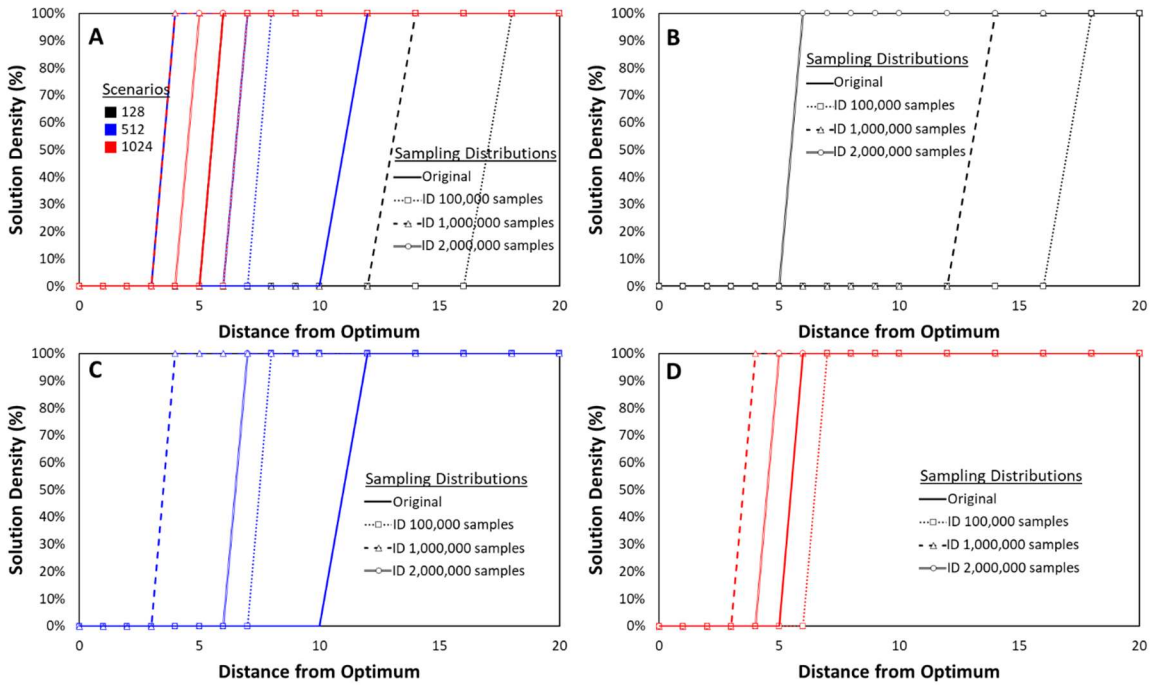

Figure 4.11 Solution density plot for Sobol sequence comparing the distributions sampled from and the scenarios generated where A) displays 128, 512, and 1,024 scenarios, B) only displays 128 scenarios, C) only displays 512 scenarios, and D) only displays 1,024 scenarios.

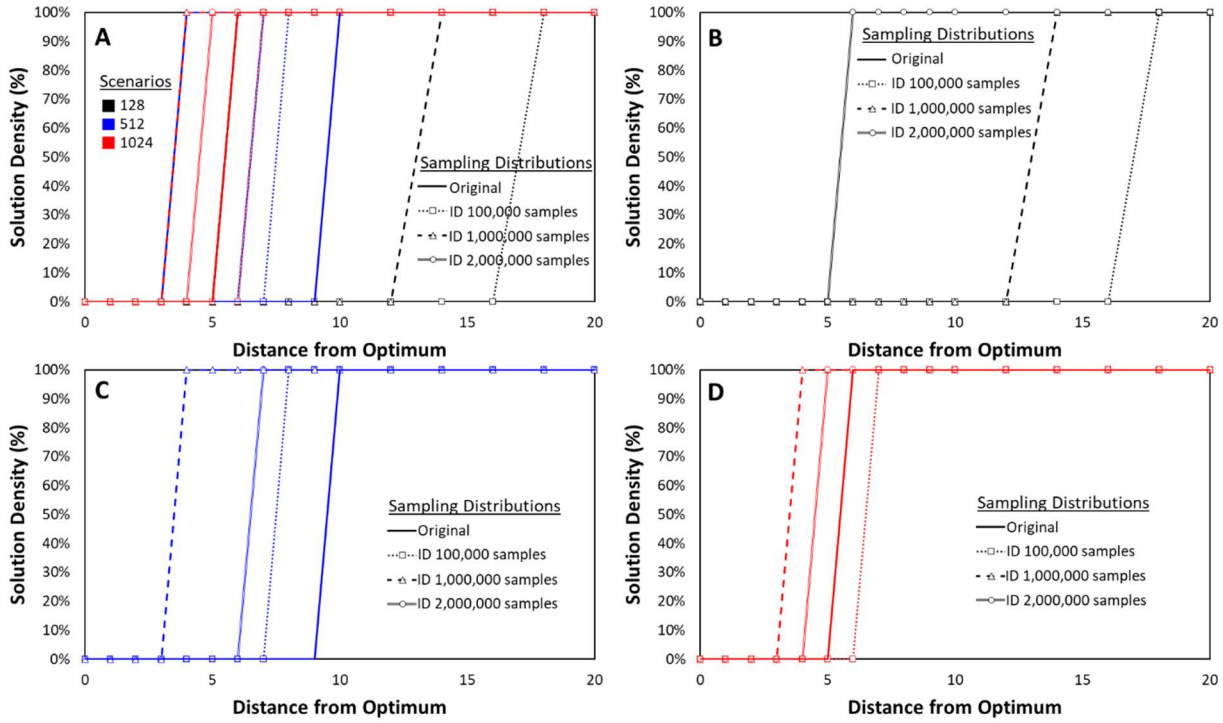

Figure 4.12 Solution density plot for Halton sequence comparing the distributions sampled from and the scenarios generated where A) displays 128, 512, and 1,024 scenarios, B) only displays 128 scenarios, C) only displays 512 scenarios, and D) only displays 1,024 scenarios.

4.1.2.2.2.3 Data Reduction Through Clustering

The recommended scenario size by the clustering methods was studied in addition to the distance from optimum and the percent difference. Plots showing all three metrics can be found in **Error! Reference source not found.**, where the color of the boxes corresponds to the number of data points used for the clustering algorithms. In **Error! Reference source not found.**, plot A shows the scenarios/clusters needed to best represent the data by each method, plot B presents the percent difference of the TSSP solution from the DFO best solution, and plot C depicts the distance from the optimum. As depicted in **Error! Reference source not found.**-A, the k-means algorithm recommended a relatively similar number of cluster centers, i.e., scenarios, across the number of data points used and when considering the added feature of our importance function (IF). As determined by the elbow method, the number of cluster centers for the k-means method is also fewer than the x-means algorithm. The number of clusters recommended by the x-means algorithm reached the upper limit when more than a million data points were clustered, which suggests that more than 1000 clusters, i.e., scenarios, may be necessary to represent the dataset best according to the BIC metric. However, a large scenario size would result in a large MILP problem, which may become intractable. In contrast to the k-means algorithm, the x-means algorithm recommended over 100 more clusters for 100,000 data points when IF was included. Furthermore, including IF decreased the variance of the recommended number of clusters by 27% for the x-means algorithm.

The TSSP solution quality is relatively similar among methods and the dataset sizes, assessed by the percent difference metric, see **Error! Reference source not found.**-B. Only a handful of TSSP solutions had over 10% difference. In general, as the size of the dataset used for clustering increased, the solution quality improved, yielding a lower percent difference. The variability of the percent difference metric is lower for the scenarios generated by the x-means

algorithm for smaller datasets. However, the variability is similar for scenarios constructed by the k-means and x-means algorithms when the dataset size is large. These results suggest that distance-based clustering algorithms could generate scenario sets that adequately represent the rare-event space to accurately estimate the benefits of the screening actions for the CRCSP.

The distance from the optimum, **Error! Reference source not found.**-C, also helps illustrate how well the clustering algorithms represent the rare-event space. Including IF in features or increasing the dataset size employed in clustering does not improve the distance from the optimum for the TSSP solutions obtained using the scenarios generated by the k-means algorithm. This behavior can also be seen in Figure 4.14, where the TSSP solution density obtained using the k-means constructed scenarios is plotted. There is no clear trend with the changes in dataset size or including IF in the feature set for the k-means algorithm, except for the slight shift to the left when IF was included for the 100,000 data points. Including IF yields a slight shift in the solution density for the 1,000,000 and 2,000,000 data points and a significant shift for the 100,000 data points for the x-means algorithm (Figure 4.15). Additionally, the distance from the optimal solution is smaller with increasing dataset size when scenarios are constructed by the x-means algorithm, which coincides with an increase in the recommended number of cluster centers, i.e., scenarios, needed to best represent the data. These observations suggest that determining the number of clusters based on the BIC may be better suited to represent the rare-event space in scenarios, albeit at the cost of increasing the TSSP model size.

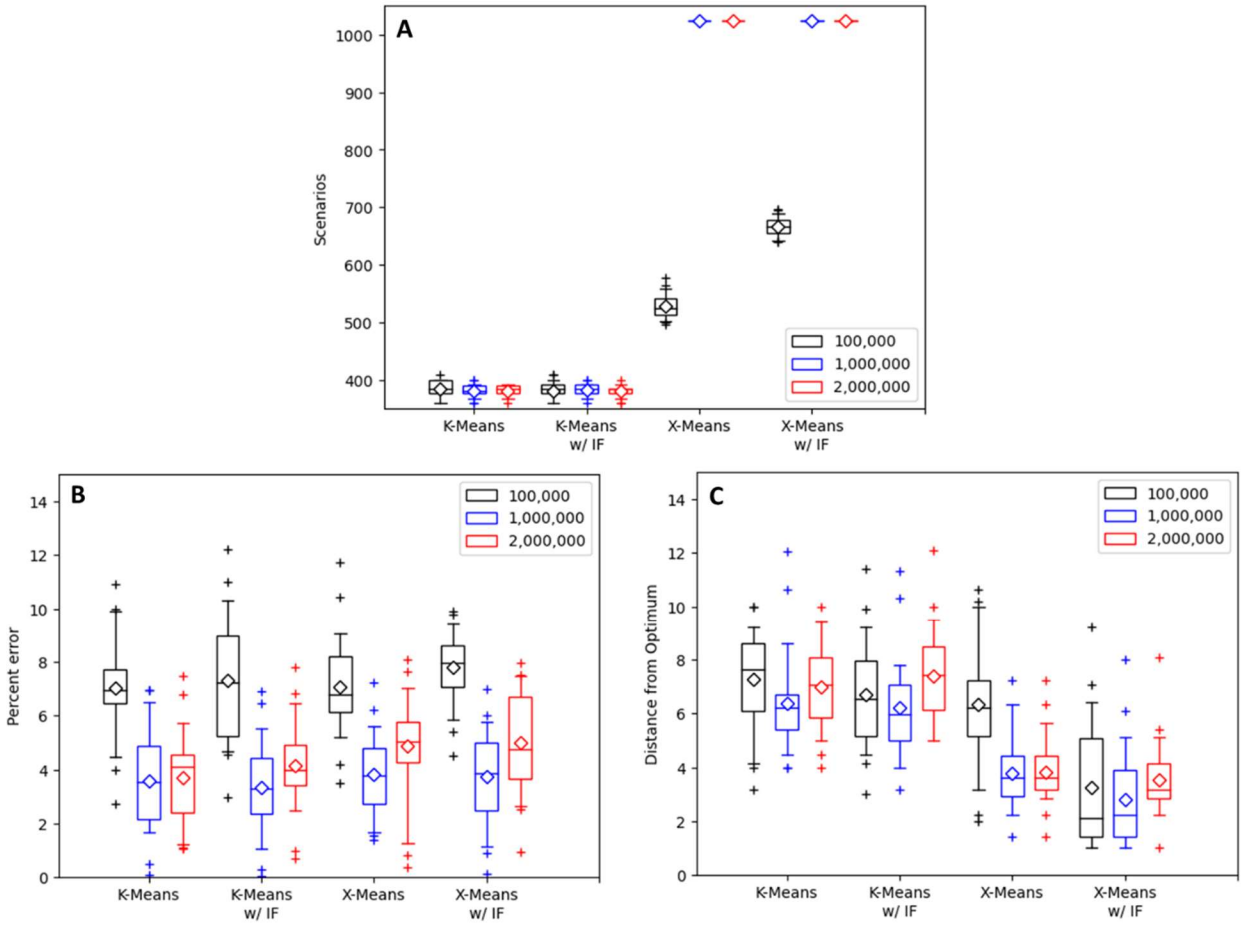

Figure 4.13 Comparison of the increase in number of data points used for clustering on the impact of the A) number of scenarios B) percent error and B) distance from the optimum when using the original features or when including the IF.

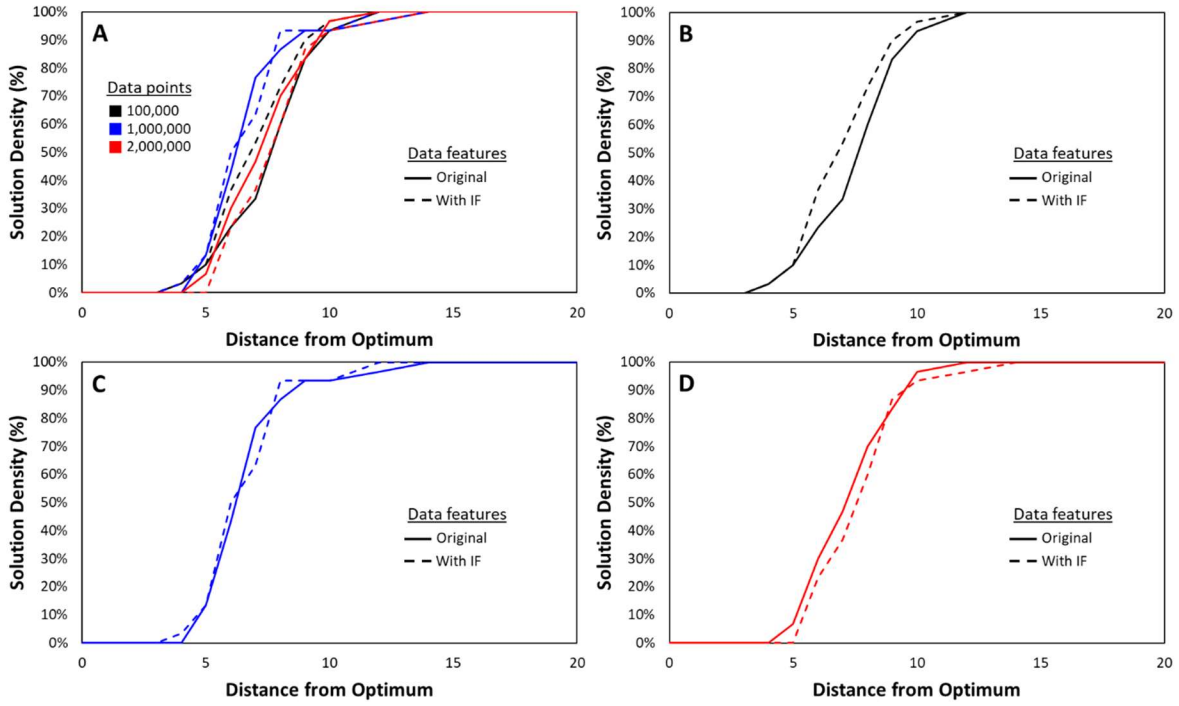

Figure 4.14 Solution density plot for k-means clustering comparing the inclusion of IF and the number of data points used where A) displays 100,000, 1,000,000, and 2,000,000 points of data, B) only displays 100,000 points of data, C) only displays 1,000,000 points of data, and D) only displays 2,000,000 points of data.

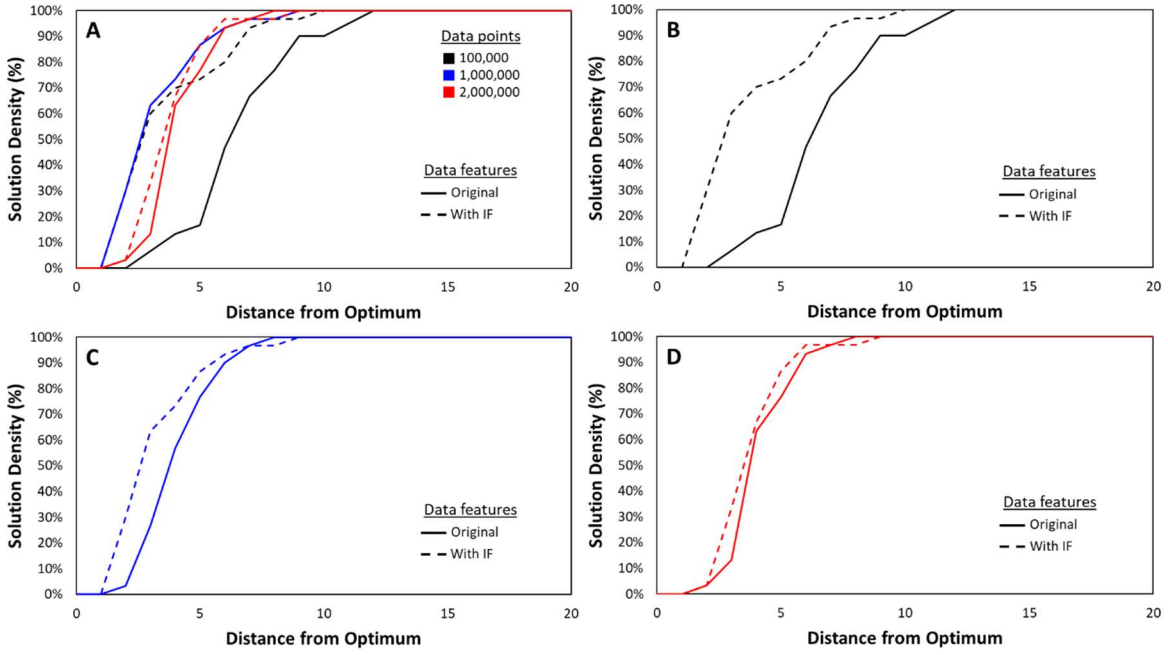

Figure 4.15 Solution density plot for x-means clustering comparing the inclusion of IF and the number of data points used where A) displays 100,000, 1,000,000, and 2,000,000 points of data, B) only displays 100,000 points of data, C) only displays 1,000,000 points of data, and D) only displays 2,000,000 points of data.

4.1.2.2.2.4 Sampling and Clustering Comparison

For this study, we evaluated ten different methodologies and their effectiveness in constructing scenarios incorporating rare-events for stochastic programs using CRCSP. We tested four different sampling methods, CMC sampling, LHS, Sobol sequence, and Halton sequence. For each sampling method, we sampled directly from the uncertain distribution and a constructed importance distribution. The last two approaches were using the k-means and x-means clustering algorithms for scenario construction. The CRCSP is a scheduling problem to determine the optimal ages to screen for colorectal cancer to maximize the gain in life for screening an asymptomatic population. It was modeled using a two-stage stochastic program whose solution heavily depends on the rare-event region of the uncertainty space, i.e., the small portion of the population that develops colorectal cancer.

We compared the performance of the methods across the solution quality as well as the number of scenarios needed, as a proxy measurement for solution time, to evaluate the overall effectiveness of the scenario construction methodologies. Figure 4.16 plots the average absolute percent difference, the distance from the optimum (d_{opt}), and number of scenarios based on if the methodology was through scenario generation, sampling, or scenario reduction methods, clustering, where each point is the average value for a given methodology, e.g., CMC sampling from the original distribution for 128 scenarios. In Figure 4.16-A, we compare the two methods based on percent difference as a function of the number of scenarios. Figure 4.16-B shows the relationship of d_{opt} and the number of scenarios, and Figure 4.16-C the comparison of the percent difference to the d_{opt} values.

In general, for the sampling methodologies, we see an expected decrease in the average value, and the spread of percent difference as the number of samples used increases. A similar

trend can be seen for the sampling methods when assessing the d_{opt} values. However, the clustering methods produce solutions with a much lower percent difference for the same or lower number of scenarios than sampling. Clustering and sampling methods construct scenarios that yield similar d_{opt} values for relatively similar scenario sizes. At smaller scenario sizes, some average values of percent difference and d_{opt} obtained using sampling methods outperform the clustering values. These instances, however, are better than clustering only on one of the metrics. As seen in plot Figure 4.16-C, the TSSP solutions obtained using the scenarios generated by sampling methodologies and have low percent differences have a higher d_{opt} value and vice versa. However, the solutions obtained with the scenarios constructed by the clustering methodologies provide both low percent differences and d_{opt} values.

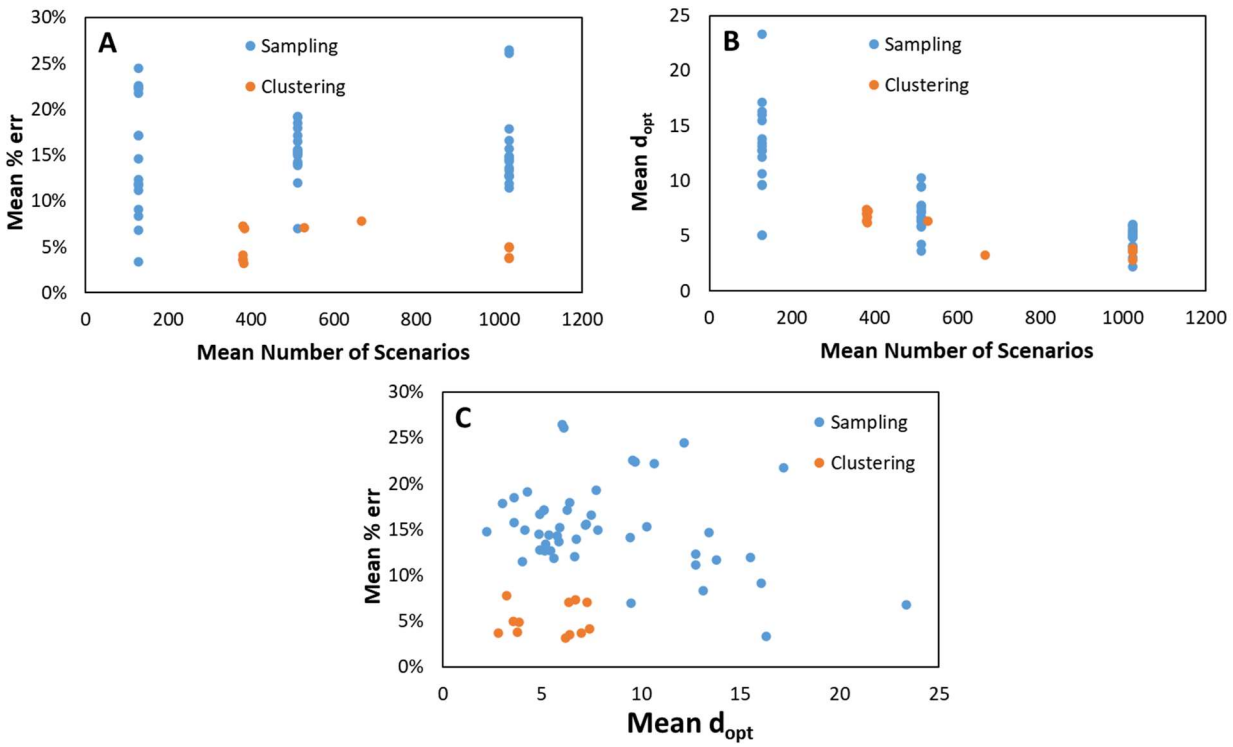

Figure 4.16 Overall comparison of scenario construction methodology comparing A) average % $\boldsymbol{err}$ versus average number of scenarios, B) average $\boldsymbol{d_{opt}}$ versus average number of scenarios, and C) average % $\boldsymbol{err}$ versus average $\boldsymbol{d_{opt}}$.

4.2 Multi-stage Stochastic Programming Approach to Solve CRCSP

In practice, it is common for health care providers to recommend when the next screening test should occur based on the results of the current screening test. This type of screening strategy is called a dynamic or observation-based screening strategy. Several studies concluded that observation-based screening strategies result in higher screening benefits than those with predetermined screening ages[68,69,76]. Stochastic programming can be used to model problems with sequential decisions as uncertainty is gradually realized. Such problems must be modeled as a multi-stage SP (MSSP), with each stage representing a decision point.

A general formulation of a MSSP can be found in Eqns. 4.67-82. The objective function is defined by Eqn. 4.67, where $\boldsymbol{x}_n$ are the decision variables associated with stage n of the problem, $\boldsymbol{\omega}_n$ is the uncertainty vector realized at stage n, and $Q_n(\boldsymbol{x}_{n-1}, \boldsymbol{\omega}_n)$ is the optimum utility gained for stage n due to decisions $\boldsymbol{x}_{n-1}$ and uncertainty realization $\boldsymbol{\omega}_n$. Eqn. 4.68 defines the constraints of the problem for the n-th stage of the problem, whereas Eqn. 4.69 defines the constraints for recourse actions of stage n, given decision variables $\boldsymbol{x}_{n-1}$ and realized uncertainty $\boldsymbol{\omega}_n$.

$$\min z = \min f(\boldsymbol{x}_1) + \mathbb{E}\big[Q_2(\boldsymbol{x}_1, \boldsymbol{\omega}_2) + \mathbb{E}[Q_3(\boldsymbol{x}_2, \boldsymbol{\omega}_3) + \cdots + \mathbb{E}[Q_N(\boldsymbol{x}_{N-1}, \boldsymbol{\omega}_N)] \dots]\big] \quad 4.67$$

$$s.t. \quad g_n(\boldsymbol{x}_n) \leq 0 \qquad \forall n \in \{1,2, \dots, N-1\} \quad 4.68$$

$$h_n(\boldsymbol{x}_{n-1}, \boldsymbol{\omega}_n) \leq 0 \qquad \forall n \in \{2,3, \dots, N\} \quad 4.69$$

The problem can then be reformulated into the more computationally tractable deterministic equivalent, defined in Eqns. 4.70-4.72, through the introduction of a vector of realizations, scenarios, of $\boldsymbol{\omega}_n$ that are indexed by set S and the probability measure for the realization of $\boldsymbol{\omega}_n^s$, $\mathbb{P}(\boldsymbol{\omega}_n^s)$.

$$\min z = \min f(\boldsymbol{x}_1) + \sum_{s\in S}\left(\prod_{n=2}^{N}\mathbb{P}(\boldsymbol{\omega}_n^s)\right)\sum_{n=2}^{N} Q_n(\boldsymbol{x}_{n-1}, \boldsymbol{\omega}_n^s) \quad 4.70$$

$$s.t. \quad g_n(\boldsymbol{x}_n) \leq 0 \qquad \forall n \in \{1,2,\dots,N-1\} \quad 4.71$$

$$h_n(\boldsymbol{x}_{n-1}, \boldsymbol{\omega}_n^s) \leq 0 \qquad \forall n \in \{2,3,\dots,N\}, \forall s \in S \quad 4.72$$

4.2.1 Problem Description

The objective CRCSP is to maximize the expected benefit gained by scheduling screening actions for an asymptomatic population with a risk of developing colorectal cancer. The disease progression is a known process, but the progression rate is unknown. The screening actions taken can either result in a detection of an unhealthy state or a healthy state. The actions taken are imperfect tests, where the probability associated with detecting an unhealthy state depends on the current health state of an individual. Given the results of the action, another screening test may be scheduled at a later date. The benefit of screening is evaluated by introducing a quality of life adjustment to each year lived, quality-adjusted life year, where a year lived in a more severe health state of CRC is worth less than a less severe state. If a screening action detects a precursor to CRC, an adenoma, it is removed, and the individual is returned to a healthy state. If the screening detects CRC, that individual then transitions to a clinically detected CRC state corresponding to the stage of CRC detected and is no longer eligible for screening. The survival years with CRC are then modified to fit the survival probability for the CRC stage detected and the age at which CRC was diagnosed[8] for that individual. The only screening action considered is the use of a colonoscopy to simplify the CRCSP model and as it is the most accurate of the tests for CRC identification[3].

4.2.2 Scenario Determination and Model Uncertainty

The intelligent representation of all potential realizations of uncertainty within a SP is necessary to maintain computational tractability. In previous work, Section 4.1, the model

uncertainty was defined as 1) the time spent in each health state of the adenoma-carcinoma sequence (Figure 2.1) and 2) the probability of each screening test detecting an abnormal health state, yielding scenarios defined as a lifetime for an individual and a given outcome of the screening decision made. This definition allowed easy tracking of screening benefits as unique lifetimes are described per scenario. Furthermore, it resulted in the association of type I endogenous uncertainty with the probability of detecting an abnormal health state. However, using this definition requires a large number of unique realizations (i.e., individuals) to represent the progression of the disease in the population (Section 4.1.2). This significantly impacts the computational tractability of a SP model as the number of scenarios are a function of both the unique realization of disease progression and the number of potential outcomes of a screening test. Additionally, this definition simplifies the disease progression where only the first instance of each new health state is considered, and the development of new adenomas is not considered. To address these shortcomings, we define the uncertainty of the CRC progression by the percentage of a population within a health state at a given age for the MSSP model. This uncertainty definition makes scenarios the screening outcomes given the number of screening actions, or stages, considered in the problem, drastically reducing the number of scenarios needed to represent the CRC progression.

We define $\boldsymbol{\omega}$ as the vector of uncertain parameters. The vector contains two components, $\boldsymbol{h}$, an $|N|-1$ dimensional vector of health state realizations, where $|N|-1$ represents the decision stages of the system, and $\boldsymbol{\psi}$, the binary outcome of each action taken on the system, representing if an unhealthy state is observed $(\psi = 1)$ or not $(\psi = 0)$. Different health states, h_n, are defined by set H, where $H = \{1: Healthy, 2: Adenoma\ present, 3: Undetected\ CRC, 4: Detected\ CRC\}$. We then define set

O as the Cartesian product $H \times [0,1]$, resulting in eight potential outcomes for each screening action taken. To further simplify the scenario tree for this first iteration of the MSSP, we assume that there are no false positive results eliminating the outcome ($\psi = 1, h = 1$) and if CRC symptoms appear there will not be a false negative result eliminating ($\psi = 0, h = 4$). These assumptions reduced the scenario outcomes to the six outlined in Table 4.6.

The set of all scenarios, set S, is defined as the Cartesian product of set O with itself $|N| - 1$ times, $S = O^{|N|-1}$, resulting in $|S| = 6^{|N|-1}$ scenarios. The scenario size can further be reduced by eliminating unrealizable outcomes and bundling certain scenario sets due to problem characteristics. As an example, let's consider a three-stage problem, i.e., $|N|$=3. The original scenario size is $|S| = 6^{|N|-1} = 6^2 = 36$. However, if the outcome of the first action is that CRC was detected symptomatically prior to the action, $o = 5$, there is no reason to consider the possibility of the individual being in any health state other than clinically detected CRC, $h = 4$. Therefore, all scenarios where the outcome of the first action is $o = 5$ can be bundled into a single scenario $(5, -)$ without loss of information. Similarly, scenarios with outcome $o = 4$, clinically detected CRC, as the first outcome are bunded into a single scenario. Additionally, the progression of health states in the adenoma-carcinoma sequence only increases in severity unless an outside action is taken. Therefore, scenarios where unhealthy states were not detected, $o = 1$ or $o = 2$, cannot have outcomes where a less severe health state is detected in the subsequent screening action, i.e., scenarios $s = (1,0)$, $s = (2,0)$, $s = (2,1)$, and $s = (2,3)$ are not realizable. Incorporating these changes reduces the number of scenarios to 22.

Table 4.6 Possible outcomes for a screening action

Outcome (o)	Parameter realizations (ψ, h)	Outcome description
0	(0,1)	Detected as healthy and truly healthy
1	(0,2)	Detected as healthy but an adenoma has developed
2	(0,3)	Detected as healthy but CRC has developed
3	(1,2)	Detected as unhealthy with an adenoma developed
4	(1,3)	Detected as unhealthy with CRC developed
5	(1,4)	Detected as unhealthy prior to screen due to CRC symptoms

The probability measure for each scenario is expressed in Eqn. 4.73, where x_{0t_0} is a binary variable representing the first stage screening decision at time point t_0 taking a value of one if the first screening action occurs at time t_0 and $x^s_{nt_n}$ is the binary variable for screening decisions made for stage n at time point t_n for scenario s. The measure $\mathbb{P}(\boldsymbol{\psi}^s|\boldsymbol{h}^s)$ is a parameter for the MSSP model because the vector of health states visited for each scenario, $\boldsymbol{h}^s$, and the outcomes of the actions taken, $\boldsymbol{\psi}^s$, are fully described for each scenario. However, the measure for the health states visited for each stage, $\mathbb{P}\left(h^s_n \middle| o^s_{n-1}, x^s_{nt_n}, \dots, x^s_{1t_1}, x_{0t_0}\right)$, depends on the timing of the actions taken and the outcomes of the previous actions, making the probability measure a decision variable and resulting in a type-1 endogenous uncertainty.

$$\begin{aligned}\mathbb{P}(\boldsymbol{\omega}_s, \boldsymbol{x}_0, \boldsymbol{x^s}) &= \prod_{\boldsymbol{n=1}}^{\boldsymbol{|N|-2}} \mathbb{P}(\psi^s_n|h^s_n)\mathbb{P}\left(h^s_n, x^s_{nt_n} \middle| o^s_{n-1}, \dots, o^s_1, x^s_{(n-1)t_{n-1}}, \dots, x^s_{1t_1}, x_{0t_0}\right) \\ &= \mathbb{P}(\boldsymbol{\psi}^s|\boldsymbol{h}^s) \prod_{n=1}^{N-1} \mathbb{P}\left(h^s_n, x^s_{nt_n} \middle| o^s_{n-1}, \dots, o^s_1, x^s_{(n-1)t_{n-1}}, \dots, x^s_{1t_1}, x_{0t_0}\right)\end{aligned} \quad 4.73$$

Using the definitions of the decision variables x_{0t} and x^s_{nt}, Eqn. 4.74 calculates $\mathbb{P}\left(h^s_n \middle| o^s_{n-1}, x^s_{nt_n}, \dots, x^s_{1t_1}, x_{0t_0}\right)$ by introducing parameters $\Pr(h|o_n, \dots, o_1, t_n, \dots, t_0)$ to define the

probability of being in health state h_n, given the outcome history $[o^s_{n-1}, \dots, o^s_1]$, and the timing of the actions taken $[t_n, \dots, t_0]$.

$$\mathbb{P}\left(h^s_n, x^s_{nt_n} \middle| o^s_{n-1}, \dots, o^s_1, x^s_{(n-1)t_{n-1}}, \dots, x^s_{1t_1}, x_{0t_0}\right) =$$

$$\sum_{t_n \in T} \dots \sum_{t_0 \in T} \Pr(h^s_n, t_n | o^s_{n-1}, \dots, o^s_1, t_{n-1}, \dots, t_0) \left(\prod_{n'=1}^{n} x^s_{n't_{n'}} \right) x_{0t_0} \quad 4.74$$

The scenario probabilities are then calculated by Eqn. 4.75.

$$\mathbb{P}(\boldsymbol{\omega}_s, \boldsymbol{x}_0, \boldsymbol{x^s}) =$$

$$\mathbb{P}(\boldsymbol{\psi}^s | \boldsymbol{h}^s) \prod_{n=1}^{|N|-2} \sum_{t_n \in T} \dots \sum_{t_0 \in T} \Pr(h^s_n, t_n | o^s_{n-1}, \dots, o^s_1, t_{n-1}, \dots, t_0) \left(\prod_{n'=1}^{n} x^s_{n't_{n'}} \right) x_{0t_0} \quad 4.75$$

Equation 4.75, however, introduces a large degree of nonlinearity into the model as the number of stages increases. Because the decision variables x_{0t} and x^s_{nt} are binary, Eqn. 4.76 always holds, where a is an arbitrary integer value.

$$(x_{0t})^a = x_{0t} \quad 4.76$$

Using Eqn. 4.76, the nonlinearities in the probability measure calculation is reduced to Eqn. 4.77.

$$\mathbb{P}(\boldsymbol{\omega}_s, \boldsymbol{x}_0, \boldsymbol{x^s}) =$$

$$\mathbb{P}(\boldsymbol{\psi}^s|\boldsymbol{h}^s) \sum_{t_{|N|-2} \in T} \dots \sum_{t_0 \in T} \left(\prod_{n=1}^{|N|-2} \Pr(h_n^s, t_n | o_{n-1}^s, \dots, o_1^s, t_{n-1}, \dots, t_0) x_{nt_n}^s \right) \Pr(h_0^s, t_0) x_{0t_0} \quad 4.77$$

4.2.3 Objective Function of the CRCSP MSSP Model

The objective of the CRCSP is to maximize the screening benefits for the population. The uncertainty definition presented in Section 4.1 enabled directly calculating the health benefits achieved through screening. Having unique lifetimes describing the progression through the disease gave a baseline and yielded a closed form expression to calculate the life-years gained through screening. In contrast, the uncertainty definition developed for the CRCSP MSSP model only specifies the proportion of the population at a given health state and not how the population progresses through the disease. This definition, albeit leading to a compact scenario construction, does not provide a closed form expression for calculating the benefits, i.e., quality-adjusted life years gained due to screening, for the objective function.

We introduce a function, $F(\boldsymbol{x}_0, \boldsymbol{x^s}, \boldsymbol{\omega})$ To approximate the objective (Eqn. 4.67). The function estimates the benefits of screening decisions, $\boldsymbol{x}_0$ and $\boldsymbol{x^s}$, for a given uncertainty realization, ω_s. Then, the objective function is given in Eqn. 4.78.

$$\max z \approx \max \mathbb{E}_\omega[F(\boldsymbol{x}_0, \boldsymbol{x^s}, \boldsymbol{\omega})] = \max \sum_{s \in S} \mathbb{P}(\boldsymbol{\omega}_s, \boldsymbol{x}_0, \boldsymbol{x^s}) F(\boldsymbol{x}_0, \boldsymbol{x^s}, \boldsymbol{\omega}_s) \quad 4.78$$

We employ surrogate modeling[189] to obtain an expression for $F(\boldsymbol{x}_0, \boldsymbol{x^s}, \boldsymbol{\omega}_s)$. Surrogate modeling is an approach to approximate a functional relationship from a given dataset to provide a predicted output value given a set of input variables[190]. In this study, the two surrogate modeling

approaches are used to approximate $F(\boldsymbol{x}_0, \boldsymbol{x}^s, \boldsymbol{\omega}_s)$, an artificial neural network model, and a random forest model.

4.2.3.1 Using Artificial Neural Networks to Approximate Screening Benefits

Artificial neural networks (ANN) were inspired by the behavior of neurons in the brain. An ANN consists of an input and an output layer connected by a number of hidden layers[191]. Each layer contains a number of neurons. The network nodes, or neurons, have weights and biases connecting the layers, with the activation function in the hidden layer determining whether a neuron will "fire" and produce a signal. Training of an ANN refers to the process of calculating the values of the weights and biases. Equations 4.79-4.80 express the relationship between the input and output of a neuron within the ANN. In the equations, set I defines the input set, and sets J and K are the sets of the neurons within the current and next layers in the network. Here, x_{ij} is an inputs, w_{ij} and b_j are model parameters determined during training, $A(.)$ is the activation function for the neuron, variable z_{jk} is an intermediate variable and input to the activation function, and $\hat{z}_{jk}$ is the output of the neuron.

$$Z_{jk} = \sum_{i \in I} w_{ij} x_i + b_j \tag{4.79}$$

$$\hat{z}_{jk} = A(z_{jk}) \tag{4.80}$$

The ANN model developed to approximate the screening benefits for the MSSP model had a single hidden layer with the rectified linear unit (ReLU)[192] as the activation function, a linear input layer, and a single neuron in the output layer. The MSSP scenarios were incorporated into the ANN as inputs using one-hot encoding[193] by introducing $|S|$ binary variables.

The resulting formulation to approximate $F(\boldsymbol{x}_0, \boldsymbol{x^s}, \boldsymbol{\omega}_s)$ using an ANN model is given in Eqns. 4.81 to 4.91[194]. Set I includes the index for the ANN inputs, where $|I| = |N| - 1 + |S|$, and $|N|$ is the number of stages in the MSSP model. Set J is the indices of the neurons in the hidden layer (The approach used to determine the number of neurons in the hidden layer is outlined in Appendix B). Equation 4.81 is the output of the ANN and gives the estimate of $F(\boldsymbol{x}_0, \boldsymbol{x^s}, \boldsymbol{\omega}_s)$. Equations 4.82 – 4.87 calculates the ReLU activation function value for the output of the hidden layer. Equation 4.88 calculates the neuron outputs in the hidden layer prior to the activation function. Equations 4.89 – 4.91 define the ANN inputs. Equation 4.89 normalizes the values of the first stage decision variables, while Eqn. 4.90 does the same for decision variables from stages two to $|N| - 1$. Equation 4.91 defines the one-hot encoded input into the ANN for the scenario definitions, where $\theta^s_{s'}$ takes a value of one if $s' = s$ and is zero otherwise.

$$F(\boldsymbol{x}_0, \boldsymbol{x^s}, \boldsymbol{\omega}_s) = \sum_{j \in J} w_j \hat{z}_j^s + b^{out} \qquad \forall s \in S \qquad 4.81$$

$$z_j^s \leq M y_j^s \qquad \forall j \in J, \forall s \in S \qquad 4.82$$

$$z_j^s \geq -M y_j^s \qquad \forall j \in J, \forall s \in S \qquad 4.83$$

$$\hat{z}_j^s \leq M y_j^s \qquad \forall j \in J, \forall s \in S \qquad 4.84$$

$$\hat{z}_j^s \geq -M y_j^s \qquad \forall j \in J, \forall s \in S \qquad 4.85$$

$$\hat{z}_j^s \leq z_j^s + M(1 - y_j^s) \qquad \forall j \in J, \forall s \in S \qquad 4.86$$

$$\hat{z}_j^s \geq z_j^s - M\left(1 - y_j^s\right) \qquad \forall j \in J, \forall s \in S \qquad 4.87$$

$$z_j^s = \sum_{i \in I} w_{ij} \hat{x}_i^s + b_j \qquad \forall j \in J, \forall s \in S \qquad 4.88$$

$$\hat{x}_0^s = \frac{\sum_{t \in T} t x_{0t} - T_{min}}{T_{max} - T_{min}} \qquad \forall s \in S \qquad 4.89$$

$$\hat{x}_n^s = \frac{\sum_{t \in T} t x_{nt}^s - T_{min}}{T_{max} - T_{min}} \qquad \forall s \in S, \forall n \in [1, \dots, |N| - 2] \qquad 4.90$$

$$\hat{x}_n^s = \theta_{s'}^s \qquad \forall s, s' \in S, \forall n \in [|N| - 1, \dots, |N| + |S|] \qquad 4.91$$

4.2.3.2 Using Random Forests Model to Approximate Screening Benefits

Random forests (RF) is a machine learning methodology that is used for classification and regression. A random forest model generates its output by combining outcomes from a set of decision trees, called a forest. Each tree is constructed independently and employs a random vector sampled from the input for its construction[195]. The output value of a single tree in the forest for a given input is the value of the final leaf node reached, and the RF model output value is the average of the outputs for every decision tree in the forest using a bootstrap aggregation approach.

To approximate the screening benefits for the MSSP model, we constructed $|S|$ RF models, one RF model for each scenario. (We also considered using one-hot encoding of the scenarios, similar to the ANN model, for building the RF model. However, it was found that the RF model was less accurate when one-hot encoding was used. These results are presented in Appendix B.2.) A mixed integer linear formulation for incorporating a RF model into the MSSP as constraints is presented in Eqns. 4.92-4.98, where set S is the set of all scenarios in the MSSP, T^s is the set of all trees in the forest for scenario $s \in$ S, L^s is the set of leaf nodes in a tree for scenario s, and P^s is a set that defines all potential combinations of 1) the trees in the forest, t, 2) the parent node, p, 3) the left child node, l, and 4) the right child nodes, r, within the constructed forest for scenario s. The output of the RF, the average of the outputs of all trees, is presented in Eqn. 4.92, where N_{trees}^s and c_{tl}^s are parameters for the number of trees in the forest and the trained predictor value for leaf node n on tree t for scenario s, respectively. In Eqn. 4.92, y_{tl}^s is a binary variable indicating the active leaf node, a leaf node that fits all of the branching decisions within the tree, n, for tree t in scenario s. The root node for each tree is forced to be active in Eqn. 4.93, and Eqn.

4.94 enforces that only one child node for each parent node may be active if the parent node itself is active. The branching conditions to select which child node to activate for each input variable are enforced in Equations 4.95-4.98. In these expressions H_P^s is the threshold value of parent node to define the splitting decisions in the tree, and M is a sufficiently large parameter value for a big-M constraint.

$$F(\boldsymbol{x}_0, \boldsymbol{x^s}, \boldsymbol{\omega}_s) = \frac{1}{N_{trees}^s} \sum_{t \in T} \sum_{l \in L} c_{tl}^s y_{tl}^s \qquad \forall s \in S \qquad 4.92$$

$$y_{t1}^s = 1 \qquad \forall t \in T, \forall s \in S \qquad 4.93$$

$$y_{tl}^s + y_{tr}^s = y_{tp}^s \qquad \forall\ (t,p,l,r)\ \in\ P, \forall s \in S \qquad 4.94$$

$$\sum_{t \in T} tx_{0t} > H_p^s - M(1 - y_{tr}^s) - M\left(1 - y_{tp}^s\right) \qquad \forall\ (t,p,l,r)\ \in\ P, \forall s \in S \qquad 4.95$$

$$\sum_{t \in T} tx_{0t} \leq H_p^s + My_{tr}^s + M\left(1 - y_{tp}^s\right) \qquad \forall\ (t,p,l,r)\ \in\ P, \forall s \in S \qquad 4.96$$

$$\sum_{t \in T} tx_{nt}^s > H_p^s - M(1 - y_{tr}^s) - M\left(1 - y_{tp}^s\right) \qquad \forall \mathrm{n} \in [1, \dots, |N| - 2],\ \forall\ (t,p,l,r)\ \in\ P, \forall s \in S \qquad 4.97$$

$$\sum_{t \in T} tx_{nt}^s \leq H_p^s + My_{tr}^s + M\left(1 - y_{tp}^s\right) \qquad \forall \mathrm{n} \in [1, \dots, |N| - 2],\ \forall\ (t,p,l,r)\ \in\ P, \forall s \in S \qquad 4.98$$

4.2.4 Constraints and Model Formulation of the CRCSP MSSP

The rest of the MSSP model constraints incorporate the screening decisions. The first set of constraints, Eqns. 4.99-4.102, are logical constraints to define the screening actions taken. Equations 4.99 and 4.100 are used to enforce that each screening decision is taken only at one time point, and only one screening test is recommended per decision. Equations 4.101 and 4.102 are ordering constraints for the screening actions, e.g., ensuring screening action two is taken after screening action one.

$$\sum_{t\in T} x_{0t} = 1 \tag{4.99}$$

$$\sum_{t\in T} x^s_{nt} = 1 \qquad \forall \mathrm{n} \in [1, \dots, |N| - 2], \forall s \in S \tag{4.100}$$

$$\sum_{t\in T} t x_{0t} < \sum_{t\in T} t x^s_{1t} \qquad \forall s \in S \tag{4.101}$$

$$\sum_{t\in T} t x^s_{(\mathrm{n}-1)t} < \sum_{t\in T} t x^s_{nt} \qquad \forall \mathrm{n} \in [2, \dots, |N| - 2], \forall s \in S \tag{4.102}$$

Another set of constraints, non-anticipativity constraints (NACs), ensures that the future outcomes are not utilized to make here-and-now decisions. These constraints set the values of decision variables equal to each other for indistinguishable scenarios. For the CRCSP, scenarios become distinguishable based on the outcomes of the screening actions. Per each test, there are four distinguishable outcomes: 1) no unhealthy state detected, 2) an adenoma detected, 3) cancer detected, and 4) cancer detected through symptoms. In the case that no unhealthy state is detected, this could be due to either the individual being truly healthy, or the test missing an unhealthy state. However, once CRC is detected, there is no longer any reason to screen further in that scenario, meaning any outcome after CRC is detected is irrelevant and for simplicity can be bundled with any other scenario to reduce the degrees of freedom in the problem. This leads to only two outcomes needing to be differentiated from each other for the next screening decision, the no unhealthy state detected and a detected adenoma state. Meaning that number of scenario subsets needed to construct the NACs to be equal to 2^{n-1} where n is stage in the MSSP. Set $\hat{S}$ is then introduced to be the set of all subsets of scenarios described at each stage where the total number of subsets, $|\hat{S}|$, is described by $|\hat{S}| = \sum_{n=2}^{|N|-1} 2^{n-1}$. Set $\hat{S}^n \subset \hat{S}$ is then defined to be a set of all subsets needed to describe the scenario bundles at stage n. The necessary NACs for the model are then presented by Eqn. 4.103, where s_{min} is the lowest scenario index of subset J^n.

$$\sum_{t\in T} tx_{nt}^{s} = \sum_{t\in T} tx_{nt}^{s_{min}} \qquad \forall\, n \in [1,\dots,|N|-2],\ \forall s \in \hat{S}^{n}\backslash_{s_{min}} \qquad 4.103$$

Due to type I endogenous uncertainty in the CRCSP, the MSSP model is nonlinear. Following our previous approach (Section 0), the nonlinear terms are linearized by introducing new variables. The approach used is the exact linearization of products of binary variables and then products of binary and continuous variables (Section 4.1.1.2.1) using the constraints defined in Eqns. 4.104-4.110. The variable $\hat{X}^{s}_{t_0 t_1 \dots t_{|N|-2}}$ is defined to be the product of the binary variables, and $EFX^{s}_{t_0 t_1 \dots t_{|N|-2}}$ is defined to be the product of $\hat{X}^{s}_{t_0 t_1 \dots t_{|N|-2}}$ and $F(\boldsymbol{x}_0, \boldsymbol{x^s}, \boldsymbol{\omega}_s)$.

$$\sum_{t_1,\dots,t_{|N|-2}\in T^{|N|-2}} \hat{X}^{s}_{t_0 t_1 \dots t_{|N|-2}} \le x_{0t_0} \qquad \forall\, t_0 \in T, \forall s \in S \qquad 4.104$$

$$\sum_{t_{0,\dots},t_{i-1},t_{i+1},t_{|N|-2}\in T^{|N|-2}} \hat{X}^{s}_{t_0 t_1 \dots t_{|N|-1}} \le x^{s}_{it_i} \qquad \forall i \in [1,\dots,|N|-2],\ \forall\, t_i \in T, \forall s \in S \qquad 4.105$$

$$\hat{X}^{s}_{t_0 t_1 \dots t_{|N|-2}} \ge x_{0t_0} + x^{s}_{1t_1} + \cdots + x^{s}_{|N|-2} - (|N|-2) \qquad \forall\, t_0,\dots,t_{|N|-2} \in T^{|N|-1},\ \forall s \in S \qquad 4.106$$

$$EFX^{s}_{t_0 t_1 \dots t_{|N|-2}} \le F(\boldsymbol{x}_0, \boldsymbol{x^s}, \boldsymbol{\omega}_s) + M^{UB}\left(1 - \hat{X}^{s}_{t_0 t_1 \dots t_{|N|-2}}\right) \qquad \forall\, t_0,\dots,t_{|N|-2} \in T^{|N|-1},\ \forall s \in S \qquad 4.107$$

$$EFX^{s}_{t_0 t_1 \dots t_{|N|-2}} \ge F(\boldsymbol{x}_0, \boldsymbol{x^s}, \boldsymbol{\omega}_s) - M^{LB}\left(1 - \hat{X}^{s}_{t_0 t_1 \dots t_{|N|-2}}\right) \qquad \forall\, t_0,\dots,t_{|N|-2} \in T^{|N|-1},\ \forall s \in S \qquad 4.108$$

$$EFX^{s}_{t_0 t_1 \dots t_{|N|-2}} \le M^{UB}\hat{X}^{s}_{t_0 t_1 \dots t_{|N|-2}} \qquad \forall\, t_0,\dots,t_{|N|-2} \in T^{|N|-1},\ \forall s \in S \qquad 4.109$$

$$EFX^{s}_{t_0 t_1 \dots t_{|N|-2}} \ge -M^{LB}\hat{X}^{s}_{t_0 t_1 \dots t_{|N|-2}} \qquad \forall\, t_0,\dots,t_{|N|-2} \in T^{|N|-1},\ \forall s \in S \qquad 4.110$$

Equations 4.111-4.112 define the MSSP model when $F(\boldsymbol{x}_0, \boldsymbol{x}^s, \boldsymbol{\omega}_s)$ is approximated using a ANN, coined MSCRC-NN, and Eqns. 4.113-4.114 define the MSSP when $F(\boldsymbol{x}_0, \boldsymbol{x}^s, \boldsymbol{\omega}_s)$ is approximated using a RF model, referred to as MSCRC-RF.

$MSCRC - NN$

$$max \sum_{s\in S} \sum_{t_0,\dots,t_{|N|-2}\in T^{|N|-1}} \Pr(h^s_{t_0}, t_0) \dots \Pr\left(h^s_{t_{|N|-2}}, t_{|N|-2} | o^s_{|N|-3}, \dots, o^s_0, t_{|N|-3}, \dots, t_0\right) \dots \quad 4.111$$

$$EFX^s_{t_0 t_1 \dots t_{|N|-2}}$$

$$s.t. \quad \text{4.81-4.91,4.99-4.110} \quad 4.112$$

$MSCRC - RF$

$$max \sum_{s\in S} \sum_{t_0,\dots,t_{|N|-2}\in T^{|N|-1}} \Pr(h^s_{t_0}, t_0) \dots \Pr\left(h^s_{t_{|N|-2}}, t_{|N|-2} | o^s_{|N|-3}, \dots, o^s_0, t_{|N|-3}, \dots, t_0\right) \dots \quad 4.113$$

$$EFX^s_{t_0 t_1 \dots t_{|N|-2}}$$

$$s.t. \quad \text{4.92-4.110} \quad 4.114$$

4.2.5 A Case Study of CRCSP MSSP with Three Stages

4.2.5.1 Data Sources and Parameter Values

We used a three-stage MSSP to evaluate the proposed model and to compare the performances of two surrogate modeling techniques in approximating the screening benefits. The MSSP model had two screening decision epochs and 22 scenarios. The data used to train the surrogate models was gathered using the microsimulation model (MSM) described in Section 3.1. The screening section of the MSM was modified to allow only two screening actions, using a colonoscopy, to be taken for each simulated individual. We used the screening section to simulate each possible outcome of the screening actions (random variable $\boldsymbol{\psi}$), with either 100% detection rate, if the outcome is defined to detect the unhealthy state, or 0% detection rate if it is not for every individual within the simulated cohort. Two metrics were then calculated to evaluate the screening benefits per individual within the MSM. The first metric is the expected quality-adjusted life-years (QALY) lived, LYL_S, where the years of life are weighted based on the health state of the individual after screening is applied given the scenario that the individual is described by. The calculation of LYL_S is presented in Eqn. 4.115, where P^S is the set of individuals within the simulation whose life time is described by scenario s, and $T = \{1,2,\dots,100\}$ is the set of all timepoints

in the simulation. H^{Full} is the set of health states considered within the MSM (elements defined in

Table 4.7). The quality of life adjustment value, w_h, depends on health state h, and the values are defined in

Table 4.7. The binary variable y_{pth}^{scrn} indicates if individual p is in health state h at timepoint t after screening, and μ_p is the disutility due to screening for individual p that is associated with the total number of colonoscopies taken in a lifetime. The value of μ_p for a single colonoscopy is assumed to be 0.0055 years[167].

$$LYL_s = \frac{1}{|P^s|} \sum_{p \in P^s} \sum_{t \in T} \sum_{h \in H^{Full}} w_h y_{pth}^{scrn} - \mu_p \qquad 4.115$$

The second metric used to assess screening benefits is the QALY gained, LYG_s. The QALY lived is calculated for the individual with and without the screening actions, and LYG_s is calculated by Eqn. 4.116, where a superscript of $scrn$ denotes the lifetime when the individual is screened and nat denotes when the individual is not screened.

$$LYG_s = \frac{1}{|P^s|} \sum_{p \in P^s} \left(\sum_{t \in T} \sum_{h \in H^{Full}} w_h \left(y_{pth}^{scrn} - y_{pth}^{nat} \right) \right) - \left(\mu_p^{scrn} - \mu_p^{nat} \right) \qquad 4.116$$

Table 4.7 Utility weightings of each health state

Health state (h)	Parameter value w_h
1 – Healthy	1.00
2 – Adenoma	1.00
3 – Undiagnosed stage 1 CRC	0.88
4 – Undiagnosed stage 2 CRC	0.88
5 – Undiagnosed stage 3 CRC	0.88
6 – Undiagnosed stage 4 CRC	0.88
7 – Clinical stage 1 CRC	0.80

8 – Clinical stage 2 CRC	0.75
9 – Clinical stage 2 CRC	0.60
10 – Clinical stage 2 CRC	0.30
11 – Deceased	0.00

The MSM was used to generate 1,000 data points for each of the 22 scenarios, with each data point corresponding to a screening strategy defined by two ages at which a colonoscopy would occur. The 1,000 input strategies were generated using Sobol' sequence, constraining the sample to ensure the first screening age was less than the second, adhering to the constraints in the MSSP. For each screening strategy, the MSM was simulated for 400 replications, with each replication simulating the lifetime of 1,000,000 males born in 1990. The hyperparameters of the ANN (number of neurons in the hidden layer and the number of training epochs) and RF (number of trees in a forest and the maximum depth per tree) models were determined using 5-fold cross-validation[196] with 15 Monte Carlo replications, withholding 20 % of the training data to be used as testing data. These results are presented in Appendix B. Once the hyperparameters were set, as given in Table 4.8, the ANN and RF models were trained using the full data set. The ANN models were trained in Python using TensorFlow, and the RF models were trained in Python using scikit-learn.

Table 4.8 Model parameters for NN and RF approximation of the QALY lived and QALY gained metrics

Model Parameter	QALY Lived	QALY Gained
NN neurons	34	23
NN training epochs	800	700
RF Trees	75	75
RF Max tree depth	9	10

Parity plots were constructed to assess the ANN and RF models' fit to the data and are presented in Figure 4.17. Plots A and B are the parity plots of the ANN and RF approximations for the QALY lived metric, and plots C and D are the same for the QALY gained metric. In the plots, the red dots correspond to the predicted and actual values, and the black line is the parity line using the normalized scale of zero to one. The closer the dots are to the parity line, the better the model predicts the values. Based on Figure 4.17, the RF model accurately approximates the QALY gained metric as the predicted values are on a very tight spread around the parity line. The predictions of the ANN model deviate more from the parity line than those of RF models for both metrics, QALY lived and gained. The ANN models predict negative values outside the normalized output range, while all RF models' predictions are within the normalized range for both metrics. The predictions of the ANN model for QALY lived, plot A, are both under and overestimates of the true values by a maximum absolute error of 0.21, and 39% of the predictions have absolute errors greater than 0.01. The ANN model overpredicts the true values at the lower range and under-predicts them at the upper range. A similar trend, though not as drastic, can be seen for the RF predictions of the QALY lived, plot B. The percentage of predictions with an absolute error greater than 0.01 is around 4%, showing that the RF model, on average, gives a better prediction than the ANN model. For the QALY gained metric, plots C and D, we observe that the RF model better predicts the true values. Less than 1% of the RF model predictions had an absolute error greater than 0.01, while the 52% of ANN model predictions had an absolute error greater than 0.01.

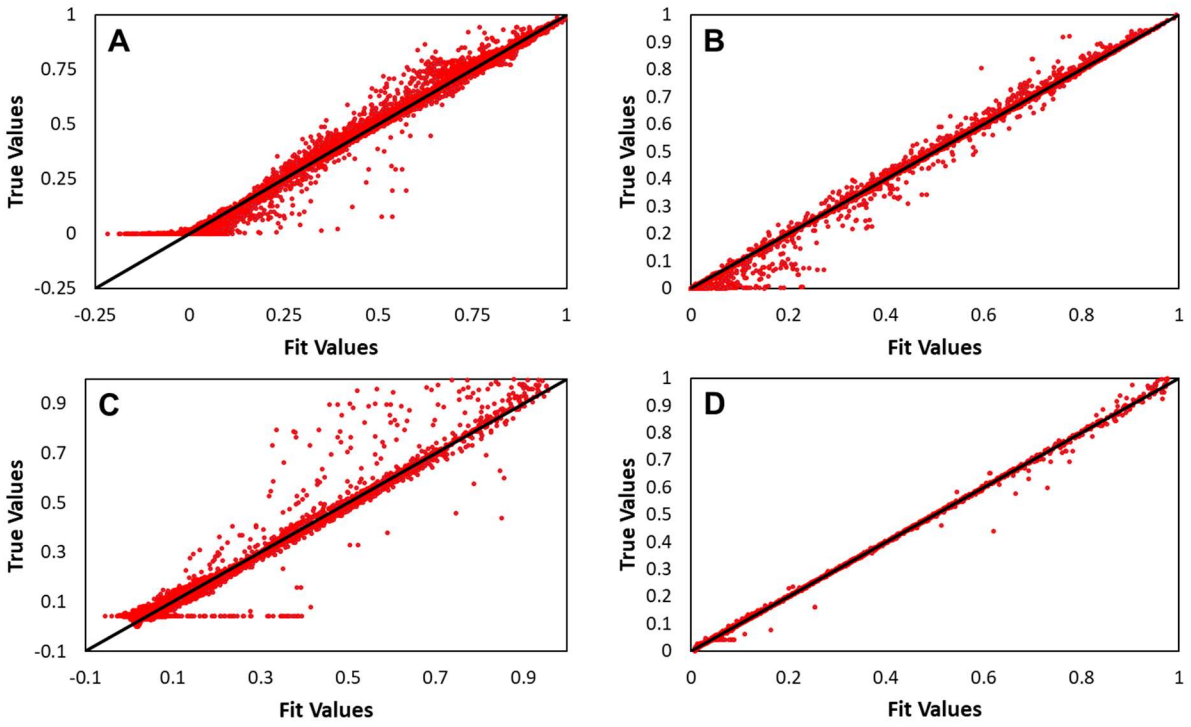

Figure 4.17 Model fit assessment of the A) NN approximation and B) RF approximation of the QALY lived metric and the C) NN approximation and D) RF approximation of the QALY gained metric.

The MSM was also used to generate data for the parameters needed to calculate the probability measures of the scenarios, $\mathbb{P}(\boldsymbol{\omega}_s, \boldsymbol{x_0}, \boldsymbol{x}^s)$. The MSM was simulated for 400 replications, with each replication simulating the lifetime of 1,000,000 males born in 1990 to determine the parameters. The probability of being at a given health state for the first stage decision, $\mathbb{P}(h_0, \mathrm{t}_0)$, was calculated by simulating a population and counting the number of individuals in each health state at every time point. For the second stage decisions, the parameter $\mathbb{P}(h_1, \mathrm{t}_1 | o_0, t_0)$ was estimated similar to the screening benefit metrics. A colonoscopy was applied at every time point, and the population was screened where the test was 100 % ($\psi = 1$) and 0 % ($\psi = 0$) accurate. After the colonoscopy, the proportion of the population was counted based on the outcome of the colonoscopy and the health state of the individuals at each time point. The

values for $\mathbb{P}(\boldsymbol{\psi}|\boldsymbol{h})$ given in Table 4.9 are defined outside of the MSM and used to represent a colonoscopy's sensitivity to a given health state.

Table 4.9 Probability of screening outcome $\mathbb{P}(\boldsymbol{\psi}|\boldsymbol{h})$

Outcome (o)	Parameter realizations $(\boldsymbol{\psi}, \boldsymbol{h})$	Probability $\mathbb{P}(\boldsymbol{\psi}|\boldsymbol{h})$
0	(0,1)	1.0
1	(0,2)	0.1
2	(0,3)	0.01
3	(1,2)	0.9
4	(1,3)	0.99
5	(1,4)	1.0

4.2.6 Model Statistics and Solutions

Both surrogate models were trained twice, once using the QALY lived (LYL) and once using the QALY gained (LYG) as outputs. The resulting model statistics are presented in Table 4.10. The MSSP model using an ANN is more compact (MSCRC-NN), with fewer binary decision variables than the MSSP model with the RF (MSCRC-RF) due to the RF model structure. In a RF model, a large number of binary variables are needed to traverse the various trees, as each branching node within a tree of a RF model introduces a binary variable. The accuracy of RF models is improved by increasing the number of trees in the model. However, in doing so, the size of the resulting MILP formulation grows. Additionally, the information passed from node to node in an ANN model is a continuous variable, unlike the series of yes-no decisions involved in a RF model. The activation function of an ANN, however, can increase the complexity of the ANN formulation by potentially introducing nonlinearities and non-convexities. The use of the ReLU activation function in this study kept the MSSP model linear.

Table 4.10 Model statistics for MSCRC-NN and MSCRC-RF using QALY lived (LYL) and QALY gained (LYG)

Model formulation	Constraints	Variables	Binary
MSCRC-NN-LYL	734,751	293,342	146,953
MSCRC-RF-LYL	1,736,611	962,8659	818,505
MSCRC-NN-LYG	731,193	292,615	146,711
MSCRC-RF-LYG	1,814,666	1,016,148	871,783

The resulting MSCRC-NN and MSCRC-RF model were constructed in Python 3.8 using Pyomo 6.1.2 and were solved to optimality using CPLEX 20.1.0 using a workstation with an Intel® Xeon® 2.30 GHz, 20-core core processor and 32 GB of RAM. The solutions and solution times are given in Table 4.11. The solution times of both formulations are comparable to, if not less than, the TSSP models of CRCL1 (Eqns. 4.59 and 4.60) and CRCL2 (Eqns. 4.61 and 4.62) of the CRCSP. These solution times are due to the reformulation of the uncertainty and the resulting compact models.

The solution of the model represents the ages at which a colonoscopy is recommended. When using the LYL as the screening benefit, both MSSP models recommend screening at a very early age, 23 and 20 for MSCRC-RF and MSCRC-NN, respectively. However, the age for the second screening was drastically different for the two MSSPs. The solution of the MSCRC-NN recommended the second screening ages of 42 and 100 for individuals with a healthy outcome and an unhealthy outcome at the first screening. In comparison, the MSCRC-RF model solution recommends the second screening at ages 69 and 70 for a healthy and unhealthy outcome. When screening benefits are evaluated using the LYL metric, the early screening age maximizes the probability of the scenarios where CRC is not developed within an individual's lifetime given the first screening test. This result suggests that the LYL is a poor metric for assessing the screening benefits, as the MSSP solution recommends employing the first screening test when a majority of

the population is healthy as to increase the overall probability of the population being healthy early in life. The age of the second screening test recommended as the solution of the MSCRC-NN-LYL model for the population with a healthy first screening outcome is consistent with the current recommendations. The second screening of the population where an unhealthy outcome occurs for the first screening test suggests that screening after detecting an adenoma is not essential to maximize the health benefits of screening, which is inherently untrue. This result further suggests that the metric of LYL is ill-suited for providing a good measure of the health benefits of screening.

The MSSP solutions when the screening utility was evaluated using the LYG metric suggest insignificant benefits. The optimum solution of the MSCRC-RF-LYG model is -3.29×10^{-3} LYG. The disutility for a colonoscopy is 5.5×10^{-3} years of life, or a total loss of QALY of 0.011 years if both screens are taken. The optimal solution suggests that the expected gain in QALY does not fully offset the screening burden. The optimal screening strategy identified as the solution of the MSCRC-NN-LYG estimates 0.77 LYG. This estimate predicts 22% more LYG than the TSSP solution, which was an overestimate. The recommended ages at which screenings should occur as determined by MSCRC-RF-LYG are 51 and 100, regardless of the outcome of the first screening test. Delaying the last screening test until the last allowable age suggests that the benefits of screening more than once in a lifetime are not worth the burden. For MSCRC-NN-LYG, the recommended screening ages are 21 and 23, given healthy and unhealthy first screening outcomes, suggesting screening later in life does not improve LYG. We hypothesize that this solution maximizes the probability of developing CRC early in life to maximize the gain in QALY for the most extreme cases.

Table 4.11 Solution statistics for MSCRC-NN and MSCRC-RF using QALY lived (LYL) and QALY gained (LYG)

Model formulation	Objective value (years)	Screen 1	Screen 2 (Healthy)	Screen 2 (Unhealthy)	Solution walltime (h:m:s)
MSCRC-NN-LYL	84.62	23	42	100	0:54:27.25
MSCRC-RF-LYL	82.32	20	69	70	2:37:54.44
MSCRC-NN-LYG	0.77	21	23	87	0:41:42.94
MSCRC-RF-LYG	-3.29×10^{-3}	51	100	100	4:46:57.45

4.3 Conclusions

This chapter outlined the formulation and evaluation of a stochastic programming approach for solving the CRCSP. Two models were investigated, a two-stage stochastic program (TSSP) and a multi-stage stochastic program (MSSP). The TSSP model assumed that only one instance of CRC or its precursor might occur in an individual. Additionally, the screening strategy determined by the TSSP does not represent the practical implementation of screening strategies by healthcare professionals, where the next screening test is recommended based on the results of the most recent one. The MSSP model rectified these two drawbacks by redefining the problem uncertainty and the structure of MSSPs modeling the sequential decision-making process. However, both problem formulations can be solved to optimality to generate solutions much faster than a simulation-optimization framework.

Three models were formulated to solve the CRCSP as a TSSP. The base model, CRCNL, a MINLP, could not be solved for any instance of the CRCSP within 24 wall hours. In comparison, the two linear models, CRCL1 and CRCL2, were solved to optimality for all but one instance of the problem. The CRCL1 model was obtained using direct linearization of the product of binary variables, and binary and continuous variables. The CRCL1 quickly became intractable due to the exponential growth in the number of variables and constraints with the increasing number of actions taken on the system. The CRCL2 model was obtained by recursive calculation of the probability-weighted utility given the actions taken on the system. This linearization approach

increased the size of the CRCL2 model at a slower rate for the increasing number of actions taken on the system than CRCL1. Although CRCL1 is a larger model than CRCL2 and grows exponentially with the number of screening tests, its solution time is 2% to 67% shorter than CRCL2's due to the tighter formulation.

The TSSP models were utilized to determine the optimal screening strategy for CRC using the simulation data for a cohort of males. A case study was devised to assess how the optimal screening strategy changed with increases in the number of screens taken in a lifetime and the impact of increasing the cohort size (which increased the number of scenarios). The results suggested that there is a diminishing benefit in the expected QALY gained for the cohort as more actions are taken within a given timeframe. The optimum solutions recommended scheduling the first screening tests to capture the majority of individuals who develop CRC. The less common CRC cases were captured with increasing the number of actions, accounting for the small increase observed in the expected QALY gained. As the cohort size increased, the time between the actions taken and the expected gain in QALY increased. This result is due to a better representation of the uncertainty space with increasing cohort size where an individual may develop CRC. The optimal actions for the smaller cohort sizes are focused on a few individuals that develop CRC in their lifetime, resulting in screening actions that are very close to each other.

The TSSP model relied on representing the rare-event region of the uncertainty space with cohort size to obtain quality solutions. Ten scenario construction approaches were evaluated to determine the most efficient method for adequately representing the rare-event region. Four different sampling methods were tested: crude Monte Carlo (CMC) sampling, Latin Hypercube sampling (LHS), Sobol sequence, and Halton sequence. For each sampling method, the scenarios were constructed by either directly sampling from the uncertain parameter distributions or a

constructed importance distribution. The last two approaches used the k-means and x-means clustering algorithms for scenario construction. The results revealed that the LHS using the importance distribution yielded scenarios that best represented the rare-event region among the sampling methods. The optimum solutions of the TSSPs using this approach had the lowest variance and were closer to the true optimum solution and the objective value. The clustering algorithms generated scenarios that yielded optimum solutions closer to the true objective function value than sampling methods. The x-means algorithm recommended a larger number of scenarios than the k-means algorithm. However, the scenarios constructed by the x-means algorithm also yielded the closest TSSP solution to the true solution. The k-means algorithm constructed scenario sets that closely approximate the objective function values with a smaller size. The scenario sets constructed by the x-means algorithm provided high-quality solutions for both objective function value and the solution location, albeit they had a larger size.

Formulating the CRCSP as an MSSP was investigated to address the assumptions and drawbacks of the TSSP model. An MSSP model is inherently a sequential decision-making approach that more closely represents the screening process implemented by health care professionals. Defining the problem uncertiainty differently in the MSSSP model drastically reduced the number of scenarios, eliminating the need for a specialized scenario construction scheme to represent the rare-event region. However, the new definition prohibited the derivation of a closed-form expression for calculating the screening benefits. Hence, surrogate models were developed to estimate the screening benefits using data generated from a CRC microsimulation model. Two machine learning techniques, ANN with a ReLU activation function and a RF model, were employed to build the surrogate models. Using either the ANN or the RD model yielded a

linear MSSP model that was solved using CPLEX. The solution times of both models on a workstation were comparable to the TSSP model solution times on a supercomputing node.

Though the MSSP model was computationally easier to solve, its solutions differ significantly from the current CRC screening recommendations. The MSSP solutions recommend screening tests on the extreme ends of the timeline, suggesting they provide minimal benefit. This result contradicts clinical findings and previous studies suggesting the metrics used to assess screening benefits or the approximations of the metrics (using the surrogate models) are not well suited for capturing the benefits past a single screening test.

CHAPTER 5

CONCLUSIONS AND FUTURE DIRECTIONS

In this dissertation, the colorectal cancer screening problem (CRCSP) was approached using two different techniques for optimization under uncertainty, simulation optimization and stochastic programming.

5.1 Simulation Optimization Approach for Solving CRCSP

In Chapter 3, a simulation optimization approach was investigated for solving the CRCSP. The investigation included the recreation and verification of the CRC-SPIN microsimulation model to simulate the progression of CRC within a population. A comparative analysis was performed to assess which derivative-free optimization (DFO) solver would perform best for combinatorial optimization problems. The results suggested that the TOMLAB/glcSolve global solver combined with the DFL local solver produced high-quality solutions with fewer function evaluations. The CRCSP was solved using the CRC-SPIN microsimulation, TOMLAB/glcSolve and DFL to determine an optimal screening strategy. The optimal screening strategy outperformed the current recommendations for CRC screening.

A sensitivity study was performed to understand the impact simulation assumptions and parameters had on the optimal screening strategy. This study showed that simulation parameters affecting the CRC progression had the greatest impact on the optimal solution. Overall, the simulation-optimization approach to solving the CRCSP allowed rigorous modeling of the CRC progression and the screening process. However, the computational cost of solving this problem was very high, and confirming global optimality is difficult. The solutions produced, however,

showed a substantial increase in the screening benefits over strategies determined through a what-if-analysis.

5.2 Mathematical Programming Approach for Solving CRCSP

Chapter 4 details the mathematical programming approach to solve the CRCSP using a stochastic programming (SP) framework with two-stage SP (TSSP) and multi-stage SP (MSSP) models. Both models incorporate type I endogenous uncertainty, where the decisions impact the probability of uncertain outcomes. Due to the uncertainty, the models were mixed-integer nonlinear programs. The TSSP model was linearized by two direct linearization procedures (CRCL1 and CRCL2), and the trade-offs of the two procedures were assessed. It was discovered that CRCL1 produced a tighter model, which grows faster than CRCL2, but had a shorter solution time.

The TSSP model of the CRCSP had a unique characteristic: the region of uncertainty that has the largest impact on the solution is a relatively small rare-event region of the space. A comparison study was conducted to determine the best method to construct a scenario set for adequately representing the rare-event space. The results suggested that the top performers were clustering algorithms, specifically k-means and x-means algorithms. The optimum solution of the TSSP model was, on average, 0.63 expected quality-adjusted life-years (QALY) gained per person.

Extending the TSSP to an MSSP model required redefining the problem uncertainty. It enabled a more realistic representation of the screening problem by allowing multiple instances of CRC to be developed within a lifetime. In the CRCSP MSSP model, a surrogate model was used to approximate a closed form expression for the screening benefits assessed using either QALY lived, or QALY gained metrics. Two machine learning models were considered, ANNs or RFs.

An instance of the model with three stages was solved, yielding an expected QALY lived of 84.62 and 82.32 years for the models using the RF and ANN as the surrogate models, respectively. An expected QALY gained of -3.29×10^{-3} and 0.77 years were obtained as the MSSP solution when the RF and ANN models were the surrogates, respectively. The resulting strategies recommended by the MSSP using both metrics suggest that the formulation does not properly represent the CRCSP.

5.3 Recommendations for Future Work

The search space in the simulation-optimization approach only considered the choice of a single screening modality when constructing the screening strategy. This search space restriction eliminates the possibility of including some currently recommended strategies, such as using a combination of a colonoscopy and FIT to screen for CRC. In future studies, redefining the decision variables to increase the screening strategy search space is recommended.

The scenario construction study only considered sampling and clustering methods for the scenario set construction. Several optimum-scenario-set construction problems using mathematical programming have been proposed in the literature. Their abilities to represent the rare-event region of an uncertainty space have not been assessed. The scenario construction study would benefit from incorporating these approaches. In addition, the scenario construction study could also investigate sampling scenario sets using a generative adversarial network (GAN).

The two metrics used to evaluate the benefit of screening for the MSSP, QALY lived and gained, did not provide a screening strategy consistent with the current recommendations. This issue may be due to the methods used to approximate the screening benefits, the representation of the scenarios, or the benefits themselves. Further investigation of how best to approximate the screening benefits through surrogate modeling is recommended. In its current form, the MSSP

implicitly handles the death of an individual by having the simulation used to generate the data handle death and incorporate it into the calculations for the screening benefits. Defining the scenarios that consider the death of an individual at the age of screening should be investigated, as it may better represent the problem at hand, potentially yielding higher quality and more accurate solutions. Finally, identifying the best metric to quantify the screening benefits and investigating the minimization of the overall cancer treatment costs, similar to the simulation-optimization objective function, may improve the MSSP model.

The MSSP formulation currently only considers the use of a colonoscopy for the screening actions and requires a large number of simulation samples if the number of stages increases. The model could be extended to allow various screening actions to be taken at each decision point. However, this extension would increase the model nonlinearity as the probabilities of detecting a non-healthy state would become a function of the decision variables. This may then require a new linearization scheme or the development of specialized solution approaches to solve the problem.

An iterative sampling scheme can be constructed to address the need for a large number of samples to construct the MSSP. A baseline set of samples could be generated in the framework, the surrogate model is trained using the samples, and the MSSP is solved. Once the MSSP is solved, the framework would construct a new set of sample points to be evaluated and included in the training set for the surrogate model. This process would be repeated until a convergence criterion is met. Additionally, the endogenous uncertainty associated with the percentage of the population at a given health state given previous screening actions could be approximated with a second surrogate model.

References

1. World Heath Organization. Cancer. (2020). Available at: https://www.who.int/news-room/fact-sheets/detail/cancer. (Accessed: 27th February 2019)
2. Knudsen, A. B. *et al.* Estimation of Benefits, Burden, and Harms of Colorectal Cancer Screening Strategies: Modeling Study for the US Preventive Services Task Force. *JAMA* **315**, 2595–2609 (2016).
3. Issa, I. A. & Noureddine, M. Colorectal cancer screening : An updated review of the available options. **23**, 5086–5096 (2017).
4. Mergener, K. & Potter, N. T. Colorectal Cancer Screening Recommendations. *Jama* **316**, 1716 (2016).
5. American Cancer Society. Colorectal Cancer Facts & Figures. (2020). Available at: https://www.cancer.org/research/cancer-facts-statistics/colorectal-cancer-facts-figures.html. (Accessed: 27th February 2019)
6. van Hees, F. *et al.* Should Colorectal Cancer Screening Be Considered in Elderly Without Prior Screening? A Cost-Effectiveness Analysis. *Ann. Intern. Med.* **160**, 750–759 (2014).
7. Bradley, C. J. *et al.* Productivity Savings from Colorectal Cancer Prevention and Control Strategies. *Am. J. Prev. Med.* **41**, e5–e14 (2011).
8. Rutter, C. M. & Savarino, J. E. An evidence-based microsimulation model for colorectal cancer: Validation and application. *Cancer Epidemiol. Biomarkers Prev.* **19**, 1992–2002 (2010).
9. Harper, P. R. & Jones, S. K. Mathematical models for the early detection and treatment of colorectal cancer. *Heal. Care Manag Sci* **8**, (2005).
10. Rose, J. *et al.* A simulation model of colorectal cancer surveillance and recurrence. *BMC*

Med. Inform. Decis. Mak. **14**, 29 (2014).

11. Parmigiani, G., Skates, S. & Zelen, M. Modeling and optimization in early detection programs with a single exam. *Biometrics* **58**, 30–36 (2002).

12. Loeve, F., Boer, R., Van Oortmarssen, G. J., Van Ballegooijen, M. & Habbema, J. D. F. The MISCAN-COLON Simulation Model for the Evaluation of Colorectal Cancer Screening. *Comput. Biomed. Res.* **32**, 13–33 (1999).

13. Orcutt, G. H. A new type of socio-economic system. *Rev. Econ. Stat.* 116–123 (1957).

14. Word, D. P., Cummings, D. A. T., Burke, D. S., Iamsirithaworn, S. & Laird, C. D. A nonlinear programming approach for estimation of transmission parameters in childhood infectious disease using a continuous time model. *J. R. Soc. Interface* **9**, 1983–1997 (2012).

15. Loeve, F. *et al.* National Polyp Study data: Evidence for regression of adenomas. *Int. J. Cancer* **111**, 633–639 (2004).

16. Soetikno, R. M. *et al.* Prevalence of Nonpolypoid (Flat and Depressed) Colorectal Neoplasms in Asymptomatic and Symptomatic Adults. *JAMA* **299**, 1027–1035 (2008).

17. Strul, H. *et al.* The prevalence rate and anatomic location of colorectal adenoma and cancer detected by colonoscopy in average-risk individuals aged 40-80 years. *Am. J. Gastroenterol.* **101**, 255–262 (2006).

18. Leshno, M., Halpern, Z. & Arber, N. Cost-effectiveness of colorectal cancer screening in the average risk population. *Health Care Manag. Sci.* **6**, 165–174 (2003).

19. Aronsson, M., Carlsson, P., Levin, L., Hager, J. & Hultcrantz, R. Cost-effectiveness of high-sensitivity faecal immunochemical test and colonoscopy screening for colorectal cancer. *Br. J. Surg.* **104**, 1078–1086 (2017).

20. Clemen, R. T. & Lacke, C. J. Analysis of Colorectal Cancer Screening Regimens. 257–267 (2001).

21. Erenay, F. S., Alagoz, O. & Said, A. Optimizing Colonoscopy Screening for Colorectal Cancer Prevention and Surveillance. *Manuf. Serv. Oper. Manag.* **16**, 381–400 (2014).

22. Gopalappa, C., Aydogan-Cremaschi, S., Das, T. K. & Orcun, S. Probability model for estimating colorectal polyp progression rates. *Health Care Manag. Sci.* **14**, 1–21 (2011).

23. Prakash, M. K. *et al.* CMOST : an open-source framework for the microsimulation of colorectal cancer screening strategies. 1–16 (2017). doi:10.1186/s12911-017-0458-9

24. Preston, A. J. & Smith, W. Disease screening designs: sensitivity and screening frequency. in *Proceedings of the annual meeting of the American statistical association* 5–9 (2001).

25. Roberts, S., Wang, L., Klein, R., Ness, R. & Dittus, R. Development of a simulation model of colorectal cancer. *ACM Trans. Model. Comput. Simul.* **18**, 1–30 (2007).

26. Frazier, A., GA, C., CS, F. & KM, K. Cost-effectiveness of screening for colorectal cancer in the general population. *JAMA* **284**, 1954–1961 (2000).

27. National Cancer Institute. Cancer Intervention and Surveillance Modeling Network (CISNET). (2020). Available at: https://cisnet.cancer.gov/. (Accessed: 27th February 2019)

28. Giovannucci, E. *et al.* Folate, methionine, and alcohol intake and risk of colorectal adenoma. *JNCI J. Natl. Cancer Inst.* **85**, 875–883 (1993).

29. Giovannucci, E. *et al.* A prospective study of cigarette smoking and risk of colorectal adenoma and colorectal cancer in US men. *JNCI J. Natl. Cancer Inst.* **86**, 183–191 (1994).

30. Giovannucci, E. *et al.* A prospective study of cigarette smoking and risk of colorectal adenoma and colorectal cancer in US women. *JNCI J. Natl. Cancer Inst.* **86**, 192–199

(1994).

31. Giovannucci, E. *et al.* Aspirin use and the risk for colorectal cancer and adenoma in male health professionals. *Ann. Intern. Med.* **121**, 241–246 (1994).

32. Giovannucci, E. *et al.* Aspirin and the risk of colorectal cancer in women. *N. Engl. J. Med.* **333**, 609–614 (1995).

33. Giovannucci, E. *et al.* Multivitamin use, folate, and colon cancer in women in the Nurses' Health Study. *Ann. Intern. Med.* **129**, 517–524 (1998).

34. Chan, A. T., Giovannucci, E. L., Schernhammer, E. S., Colditz, G. A. & others. A prospective study of aspirin use and the risk for colorectal adenoma. *Ann. Intern. Med.* **140**, 157 (2004).

35. Wei, E. K. *et al.* Comparison of risk factors for colon and rectal cancer. *Int. J. cancer* **108**, 433–442 (2004).

36. Atkin, W. S. *et al.* Once-only flexible sigmoidoscopy screening in prevention of colorectal cancer: a multicentre randomised controlled trial. *Lancet* **375**, 1624–1633 (2010).

37. DiSario, J. A., Foutch, G., Mai, H. D., Pardy, K. & Manne, R. K. Prevalence and malignant potential of colorectal polyps in asymptomatic, average-risk men. *Am. J. Gastroenterol.* **86**, (1991).

38. Neugut, A. I. & Pita, S. Role of sigmoidoscopy in screening for colorectal cancer: a critical review. *Gastroenterology* **95**, 492–499 (1988).

39. Rex, D. K., Lehman, G. A., Hawes, R. H., Ulbright, T. M. & Smith, J. J. Screening colonoscopy in asymptomatic average-risk persons with negative fecal occult blood tests. *Gastroenterology* **100**, 64–67 (1991).

40. Rickert, R. R., Auerbach, O., Garfinkel, L., Hammond, E. C. & Frasca, J. M.

Adenomatous lesions of the large bowel: an autopsy survey. *Cancer* **43**, 1847–1857 (1979).

41. Thomas, W. M., Pye, G., Hardcastle, J. D. & Walker, A. R. Screening for colorectal carcinoma: an analysis of the sensitivity of haemoccult. *Br. J. Surg.* **79**, 833–835 (1992).
42. Williams, A. R., Balasooriya, B. A. & Day, D. W. Polyps and cancer of the large bowel: a necropsy study in Liverpool. *Gut* **23**, 835–842 (1982).
43. Hardcastle, J. D. *et al.* Randomised, controlled trial of faecal occult blood screening for colorectal cancer: results for first 107 349 subjects. *Lancet* **333**, 1160–1164 (1989).
44. Hardcastle, J. D. *et al.* Randomised controlled trial of faecal-occult-blood screening for colorectal cancer. *Lancet* **348**, 1472–1477 (1996).
45. Johnson, D. A. *et al.* A prospective study of the prevalence of colonic neoplasms in asymptomatic patients with an age-related risk. *Am. J. Gastroenterol.* **85**, (1990).
46. Koretz, R. L. Malignant polyps: are they sheep in wolves' clothing? *Ann. Intern. Med.* **118**, 63–68 (1993).
47. Kronborg, O., Fenger, C., Olsen, J., Bech, K. & Søndergaard, O. Repeated screening for colorectal cancer with fecal occult blood test: a prospective randomized study at Funen, Denmark. *Scand. J. Gastroenterol.* **24**, 599–606 (1989).
48. Kronborg, O., Fenger, C., Olsen, J., Jørgensen, O. D. & Søndergaard, O. Randomised study of screening for colorectal cancer with faecal-occult-blood test. *Lancet* **348**, 1467–1471 (1996).
49. Lieberman, D. A. & Smith, F. W. Screening for colon malignancy with colonoscopy. *Am. J. Gastroenterol.* **86**, (1991).
50. Mandel, J. S. *et al.* Reducing mortality from colorectal cancer by screening for fecal

occult blood. *N. Engl. J. Med.* **328**, 1365–1371 (1993).

51. Surveillance, Epidemiology, and End Results Program. *SEER*

52. Rutter, C. M., Miglioretti, D. L. & Savarino, J. E. Bayesian Calibration of Microsimulation Models. *J Am Stat Assoc.* **104**, 1338–1350 (2010).

53. Pickhardt, P. J. *et al.* Computed Tomographic Virtual Colonoscopy to Screen for Colorectal Neoplasia in Asymptomatic Adults. *N. Engl. J. Med.* **349**, 2191–2200 (2003).

54. Lieberman, D. A. *et al.* Use of Colonoscopy to Screen Asymptomatic Adults for Colorectal Cancer. *N. Engl. J. Med.* **343**, 162–168 (2000).

55. Imperiale, T. F. *et al.* Risk of Advanced Proximal Neoplasms in Asymptomatic Adults According to the Distal Colorectal Findings. *N. Engl. J. Med.* **343**, 169–174 (2000).

56. Church, J. M. Clinical significance of small colorectal polyps. *Dis. Colon Rectum* **47**, 481–485 (2004).

57. Odom, S. R., Duffy, S. D., Barone, J. E., Vishal, G. & McClane, S. J. The rate of adenocarcinoma in endoscopically removed colorectal polyps. *Am. Surg.* **71**, 1024–1026 (2005).

58. Greuter, M. J. E. *et al.* Modeling the adenoma and serrated pathway to colorectal cancer (ASCCA). *Risk Anal.* **34**, (2014).

59. De Wijkerslooth, T. R. *et al.* Immunochemical fecal occult blood testing is equally sensitive for proximal and distal advanced neoplasia. *Am. J. Gastroenterol.* **107**, 1570 (2012).

60. Tappenden, P. *et al.* Option appraisal of population-based colorectal cancer screening programmes in England. *Gut* **56**, 677–684 (2007).

61. Telford, J. J., Levy, A. R., Sambrook, J. C., Zou, D. & Enns, R. A. The cost-effectiveness

of screening for colorectal cancer. *Cmaj* **182**, 1307–1313 (2010).

62. Goede, S. L. *et al.* Cost-effectiveness of one versus two sample faecal immunochemical testing for colorectal cancer screening. *Gut* **62**, 727–734 (2013).

63. Brenner, H., Chang-Claude, J., Seiler, C. M., Stürmer, T. & Hoffmeister, M. Does a negative screening colonoscopy ever need to be repeated? *Gut* **55**, 1145–1150 (2006).

64. Imperiale, T. F. *et al.* Five-year risk of colorectal neoplasia after negative screening colonoscopy. *N. Engl. J. Med.* **359**, 1218–1224 (2008).

65. Force, U. S. P. S. T. & others. Screening for colorectal cancer: US Preventive Services Task Force recommendation statement. *Ann. Intern. Med.* **149**, 627 (2008).

66. Knudsen, A. B. *et al.* Cost-effectiveness of computed tomographic colonography screening for colorectal cancer in the medicare population. *J. Natl. Cancer Inst.* **102**, 1238–1252 (2010).

67. Greuter, M. J. E. *et al.* Modeling and optimization in early detection programs with a single exam. *Biometrics* **58**, 30–36 (2002).

68. Song, L.-P. & Wang, H.-Y. Modeling and Control of Colorectal Cancer. *PLoS One* **11**, e0161349 (2016).

69. Li, Y., Zhu, M., Klein, R. & Kong, N. Using a partially observable Markov chain model to assess colonoscopy screening strategies - A cohort study. *Eur. J. Oper. Res.* **238**, 313–326 (2014).

70. Rauner, M. S., Gutjahr, W. J., Heidenberger, K., Wagner, J. & Pasia, J. Dynamic Policy Modeling for Chronic Diseases: Metaheuristic-Based Identification of Pareto-Optimal Screening Strategies. *Oper. Res.* **58**, 1269–1286 (2010).

71. Pignone, M. Cost-Effectiveness Analyses of Colorectal Cancer Screening. *Ann. Intern.*

Med. **137**, 96 (2002).

72. Zauber, A. G. *et al.* Cost-Effectiveness of CT Colonography to Screen for Colorectal Cancer Cost-Effectiveness of CT Colonography to Screen for Colorectal Cancer Cancer Intervention and Surveillance Modeling Network (CISNET) for MISCAN, SimCRC, and CRC-SPIN Models. 1–92 (2009).

73. Haug, U. *et al.* Development of new non-invasive tests for colorectal cancer screening: The relevance of information on adenoma detection. **136**, 2864–2874 (2016).

74. Amaran, S., Sahinidis, N. V., Sharda, B. & Bury, S. J. Simulation optimization: a review of algorithms and applications. *Ann. Oper. Res.* **240**, 351–380 (2016).

75. Rios, L. M. & Sahinidis, N. V. Derivative-free optimization: A review of algorithms and comparison of software implementations. *J. Glob. Optim.* **56**, 1247–1293 (2013).

76. McLay, L. A., Foufoulides, C. & Merrick, J. R. W. Using simulation-optimization to construct screening strategies for cervical cancer. *Health Care Manag. Sci.* **13**, 294–318 (2010).

77. Bertsimas, D., Silberholz, J. & Trikalinos, T. Optimal healthcare decision making under multiple mathematical models: application in prostate cancer screening. *Health Care Manag. Sci.* **21**, 105–118 (2018).

78. Das, S. & Suganthan, P. N. Differential evolution: A survey of the state-of-the-art. *IEEE Trans. Evol. Comput.* **15**, 4–31 (2011).

79. Martelli, E. & Amaldi, E. PGS-COM: A hybrid method for constrained non-smooth black-box optimization problems. Brief review, novel algorithm and comparative evaluation. *Comput. Chem. Eng.* **63**, 108–139 (2014).

80. Boukouvala, F., Misener, R. & Floudas, C. A. Global optimization advances in Mixed-

Integer Nonlinear Programming, MINLP, and Constrained Derivative-Free Optimization, CDFO. *Eur. J. Oper. Res.* **252**, 701–727 (2016).

81. Liuzzi, G., Lucidi, S. & Rinaldi, F. Derivative-Free Methods for Mixed-Integer Constrained Optimization Problems. *J. Optim. Theory Appl.* **164**, 933–965 (2015).
82. Audet, C., Custódio, A. L. & Dennis, J. E. Erratum: Mesh Adaptive Direct Search Algorithms for Constrained Optimization. *SIAM J. Optim.* **18**, 1501–1503 (2008).
83. Rainville, F. De, Fortin, F., Gardner, M., Parizeau, M. & Gagné, C. DEAP : A Python Framework for Evolutionary Algorithms. *Companion proc. Genet. Evol. Comput. Conf.* 85–92 (2012). doi:doi:10.1145/2330784.2330799
84. Mitchell, M. *An introduction to genetic algorithms*. (MIT press, 1998).
85. Hansen, N. The CMA Evolution Strategy: A Tutorial. (2016).
86. Hansen, N. A CMA-ES for Mixed-Integer Nonlinear Optimization. *INRIA* **[Researc R**, (2011).
87. Lacerda. Simulated annealing. (2006).
88. The GPyOpt authors. {GPyOpt}: A Bayesian Optimization framework in python. (2016).
89. Frazier, P. I. A Tutorial on Bayesian Optimization. 1–22 (2018).
90. Holmstr, K., Anders, O. G. & Edvall, M. M. User's Guide for TOMLAB 7. (2010).
91. Jones, D. R. Direct Global Optimization Algorithm. *Encyclopedia of Otimization* **1**, 725–735 (2010).
92. Schlüter, M., Egea, J. A. & Banga, J. R. Extended ant colony optimization for non-convex mixed integer nonlinear programming. *Comput. OR* **36**, 2217–2229 (2009).
93. Schlüter, M. & Gerdts, M. The oracle penalty method. *J. Glob. Optim.* **47**, 293–325 (2010).

94. Schlüter, M., Egea, J. A., Antelo, L. T., Alonso, A. A. & Banga, J. R. An extended ant colony optimization algorithm for integrated process and control system design. *Ind. Eng. Chem. Res.* **48**, 6723–6738 (2009).

95. Costa, A. & Nannicini, G. RBFOpt : an open-source library for black-box optimization with costly function evaluations. *Optim. online no.4538* (2014).

96. McKay, M. D., Beckman, R. J. & Conover, W. J. Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code. *Technometrics* **21**, 239–245 (1979).

97. Birge, J. R. & Louveaux, F. *Introduction to stochastic programming*. (Springer Science & Business Media, 2011).

98. Shapiro, A., Dentcheva, D. & Ruszczyński, A. Lectures on Stochastic Programming: Modeling and Theory, Second Edition. *Lect. Stoch. Program. Model. Theory, Second Ed.* (2014). doi:10.1137/1.9781611973433

99. Solak, S., Clarke, J.-P. B., Johnson, E. L. & Barnes, E. R. Optimization of R&D project portfolios under endogenous uncertainty. *Eur. J. Oper. Res.* **207**, 420–433 (2010).

100. Høyland, K. & Wallace, S. W. Generating scenario trees for multistage decision problems. *Manage. Sci.* **47**, 295–307 (2001).

101. Calfa, B. A., Agarwal, A., Grossmann, I. E. & Wassick, J. M. Data-driven multi-stage scenario tree generation via statistical property and distribution matching. *Comput. Chem. Eng.* **68**, 7–23 (2014).

102. Pflug, G. Scenario tree generation for multiperiod financial optimization by optimal discretization. *Math. Program.* **89**, 251–271 (2001).

103. Dupačová, J., Gröwe-Kuska, N. & Römisch, W. Scenario reduction in stochastic

programming: An approach using probabilty metrics. *Math. Program.* **95**, 493–511 (2003).

104. Heitsch, H. & Römisch, W. Scenario reduction in stochastic programming. *Comput. Optim. Appl.* **24**, 187–206 (2003).

105. Heitsch, H. & Römisch, W. A note on scenario reduction for two-stage stochastic programs. *Oper. Res. Lett.* **35**, 731–738 (2007).

106. Heitsch, H., Römische, W. & Strugarek, C. Stability of multistage stochastic programs. *SIAM J. Optim.* **17**, 511–525 (2006).

107. Heitsch, H. & Römisch, W. Scenario tree reduction for multistage stochastic programs. *Comput. Manag. Sci.* **6**, 117–133 (2009).

108. Li, Z. & Floudas, C. A. Optimal scenario reduction framework based on distance of uncertainty distribution and output performance: I. Single reduction via mixed integer linear optimization. *Comput. \& Chem. Eng.* **70**, 50–66 (2014).

109. Li, Z. & Floudas, C. A. Optimal scenario reduction framework based on distance of uncertainty distribution and output performance: II. Sequential reduction. *Comput. \& Chem. Eng.* **84**, 599–610 (2016).

110. Li, Z. & Li, Z. Linear programming-based scenario reduction using transportation distance. *Comput. \& Chem. Eng.* **88**, 50–58 (2016).

111. Karuppiah, R., Martin, M. & Grossmann, I. E. A simple heuristic for reducing the number of scenarios in two-stage stochastic programming. *Comput. Chem. Eng.* **34**, 1246–1255 (2010).

112. Park, S., Xu, Q. & Hobbs, B. F. Comparing scenario reduction methods for stochastic transmission planning. *IET Gener. Transm. Distrib.* **13**, 1005–1013 (2019).

113. Li, J., Zhou, J. & Chen, B. Review of wind power scenario generation methods for optimal operation of renewable energy systems. *Appl. Energy* **280**, 115992 (2020).

114. Seljom, P., Kvalbein, L., Hellemo, L., Kaut, M. & Ortiz, M. M. Stochastic modelling of variable renewables in long-term energy models: Dataset, scenario generation & quality of results. *Energy* **236**, (2021).

115. Ehrenstein, M., Wang, C.-H. & Guillén-Gosálbez, G. Strategic planning of supply chains considering extreme events: Novel heuristic and application to the petrochemical industry. *Comput. Chem. Eng.* **125**, 306–323 (2019).

116. Garcia-Herreros, P., Wassick, J. M. & Grossmann, I. E. Design of resilient supply chains with risk of facility disruptions. *Ind. Eng. Chem. Res.* **53**, 17240–17251 (2014).

117. Feng, Y. & Ryan, S. M. Solution sensitivity-based scenario reduction for stochastic unit commitment. *Comput. Manag. Sci.* **13**, 29–62 (2016).

118. Jonsbråten, T. W. Optimization models for petroleum field exploitation. (1998).

119. Goel, V. & Grossmann, I. E. A Class of stochastic programs with decision dependent uncertainty. *Math. Program.* **108**, 355–394 (2006).

120. Tarhan, B. & Grossmann, I. E. A multistage stochastic programming approach with strategies for uncertainty reduction in the synthesis of process networks with uncertain yields. *Comput. Chem. Eng.* **32**, 766–788 (2008).

121. Dupačová, J. Optimization under exogenous and endogenous uncertainty. *Proc. Math. Methods Eng. Int. Symp.* (2006).

122. Hellemo, L., Barton, P. I. & Tomasgard, A. Decision-dependent probabilities in stochastic programs with recourse. *Comput. Manag. Sci.* **15**, 369–395 (2018).

123. Pflug, G. C. On-line optimization of simulated Markovian processes. *Math. Oper. Res.* **15**,

381–395 (1990).

124. Jonsbråten, T. W., Wets, R. J. B. & Woodruff, D. L. A class of stochastic programs withdecision dependent random elements. *Ann. Oper. Res.* **82**, 83–106 (1998).

125. Ahmed, S. Strategic planning under uncertainty: Stochastic integer programming approaches. *ProQuest Diss. Theses* 145 (2000).

126. Peeta, S., Sibel Salman, F., Gunnec, D. & Viswanath, K. Pre-disaster investment decisions for strengthening a highway network. *Comput. Oper. Res.* **37**, 1708–1719 (2010).

127. Du, L. & Peeta, S. A Stochastic Optimization Model to Reduce Expected Post-Disaster Response Time Through Pre-Disaster Investment Decisions. *Networks Spat. Econ.* **14**, 271–295 (2014).

128. Medal, H. R., Pohl, E. A. & Rossetti, M. D. Allocating Protection Resources to Facilities When the Effect of Protection is Uncertain. *IIE Trans.* **48**, 220–234 (2016).

129. Bhuiyan, T. H., Moseley, M. C., Medal, H. R., Rashidi, E. & Grala, R. K. A stochastic programming model with endogenous uncertainty for incentivizing fuel reduction treatment under uncertain landowner behavior. *Eur. J. Oper. Res.* **277**, 699–718 (2019).

130. Krasko, V. & Rebennack, S. Two-stage stochastic mixed-integer nonlinear programming model for post-wildfire debris flow hazard management: Mitigation and emergency evacuation. *Eur. J. Oper. Res.* **263**, 265–282 (2017).

131. Laumanns, M., Prestwich, S. & Kawas, B. Distribution shaping and scenario bundling for stochastic programs with endogenous uncertainty ∗. *EURO-INFORMS Jt. Int. Meet. 2013 Rome Int. Conf. Oper. Res. (OR 2014) Aachen* 1–24 (2014).

132. Escudero, L. F., Garín, M. A., Monge, J. F. & Unzueta, A. On preparedness resource allocation planning for natural disaster relief under endogenous uncertainty with time-

consistent risk-averse management. *Comput. Oper. Res.* **98**, 84–102 (2018).

133. da Costa Flach, B. Stochastic programming with endogenous uncertainty: an application in humanitarian logistics. (PhD thesis, PUC-Rio, 2010).

134. Bhuiyan, T. H., Medal, H. R. & Harun, S. A stochastic programming model with endogenous and exogenous uncertainty for reliable network design under random disruption. *Eur. J. Oper. Res.* **285**, 670–694 (2020).

135. Van Slyke, R. M. & Wets, R. L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM J. Appl. Math.* **17**, 638–663 (1969).

136. Tawarmalani, M. & Sahinidis, N. V. A polyhedral branch-and-cut approach to global optimization. *Math. Program.* **103**, 225–249 (2005).

137. Ekin, T. Integrated maintenance and production planning with endogenous uncertain yield. *Reliab. Eng. Syst. Saf.* **179**, 52–61 (2018).

138. Leo, E. & Engell, S. Condition-based maintenance optimization via stochastic programming with endogenous uncertainty. *Comput. Chem. Eng.* **156**, 107550 (2022).

139. Yin, W., Peng, X. & Hou, Y. A decision-dependent stochastic approach for wind farm maintenance scheduling considering wake effect. *IEEE PES Innov. Smart Grid Technol. Conf. Eur.* **2020-Octob**, 814–818 (2020).

140. Ekin, T., Polson, N. G. & Soyer, R. Augmented nested sampling for stochastic programs with recourse and endogenous uncertainty. *Nav. Res. Logist.* **64**, 613–627 (2017).

141. Floudas, C. A. & Visweswaran, V. A global optimization algorithm (GOP) for certain classes of nonconvex NLPs—I. Theory. *Comput. \& Chem. Eng.* **14**, 1397–1417 (1990).

142. Hu, B. *et al.* Decision-Dependent Uncertainty Modeling in Power System Operational Reliability Evaluations. *IEEE Trans. Power Syst.* **36**, 5708–5721 (2021).

143. Zhan, Y., Zheng, Q. P., Wang, J. & Pinson, P. Generation Expansion Planning with Large Amounts of Wind Power via Decision-Dependent Stochastic Programming. *IEEE Trans. Power Syst.* **32**, 3015–3026 (2017).

144. Ma, S., Su, L., Wang, Z., Qiu, F. & Guo, G. Resilience enhancement of distribution grids against extreme weather events. *IEEE Trans. Power Syst.* **33**, 4842–4853 (2018).

145. Tong, K., Feng, Y. & Rong, G. Planning under demand and yield uncertainties in an oil supply chain. *Ind. Eng. Chem. Res.* **51**, 814–834 (2012).

146. Chen, Y., Ye, Y., Yuan, Z., Grossmann, I. E. & Chen, B. Integrating stochastic programming and reliability in the optimal synthesis of chemical processes. *Comput. Chem. Eng.* **157**, 107616 (2022).

147. Bhuiyan, T. H., Moseley, M. C., Medal, H. R., Rashidi, E. & Grala, R. K. A stochastic programming model with endogenous uncertainty for incentivizing fuel reduction treatment under uncertain landowner behavior. *Eur. J. Oper. Res.* **277**, 699–718 (2019).

148. Zhang, Z., Denton, B. T. & Morgan, T. M. Optimization of Active Surveillance Strategies for Heterogeneous Patients with Prostate Cancer. *Ssrn* 1–36 (2019). doi:10.2139/ssrn.3486629

149. Türkay, M. & Grossmann, I. E. Logic-based MINLP algorithms for the optimal synthesis of process networks. *Comput. \& Chem. Eng.* **20**, 959–978 (1996).

150. Chen, Q. & Grossmann, I. Modern modeling paradigms using generalized disjunctive programming. *Processes* **7**, 839 (2019).

151. Alvarado, M. & Ntaimo, L. Chemotherapy appointment scheduling under uncertainty using mean-risk stochastic integer programming. *Health Care Manag. Sci.* **21**, 87–104 (2018).

152. MADHAVI, K., TIRUPATHIRAO, P., REDDY, P. R. S. Optimal Drug Administration for Cancer. *Am. J. Appl. Math. Math. Sci.* **2**, 37–45 (2013).

153. Sir, M. Y., Epelman, M. A. & Pollock, S. M. Stochastic programming for off-line adaptive radiotherapy. *Ann. Oper. Res.* **196**, 767–797 (2012).

154. Zaghian, M., Lim, G. J. & Khabazian, A. A chance-constrained programming framework to handle uncertainties in radiation therapy treatment planning. **266**, 736–745 (2018).

155. Vieira, B., Demirtas, D., van de Kamer, J. B., Hans, E. W. & van Harten, W. A mathematical programming model for optimizing the staff allocation in radiotherapy under uncertain demand. *Eur. J. Oper. Res.* **270**, 709–722 (2018).

156. Inadomi, J. M. *et al.* Adherence to Colorectal Cancer Screening. *Arch. Intern. Med.* **172**, 575 (2012).

157. Rutter, C. M. & Savarino, J. E. An evidence-based microsimulation model for colorectal cancer: validation and application. *Cancer Epidemiol. Prev. Biomarkers* **19**, 1992–2002 (2010).

158. Ries, L. A. G., Kosary, C. L., Hankey, B. F., Miller, B. A. & Edwards, B. K. SEER cancer statistics review, 1973-1995. *Bethesda, MD Natl. Cancer Inst.* (1998).

159. Czepiel, S. A. Maximum Likelihood Estimation of Logistic Regression Models : Theory and Implementation.

160. Cox, D. R. Regression Models and Life-Tables. *J. R. Stat. Soc. Ser. B* **34**, 187–220 (1972).

161. Yu, B., Tiwari, R. C., Cronin, K. A., Mcdonald, C. & Feuer, E. J. CANSURV : A Windows program for population-based cancer survival analysis. (2005). doi:10.1016/j.cmpb.2005.08.002

162. Kuntz, K. M. *et al.* A systematic comparison of microsimulation models of colorectal

cancer: the role of assumptions about adenoma progression. *Med. Decis. Making* **31**, 530–9 (2011).

163. Winawer, S. J. *et al.* Guidelines for Colonoscopy Surveillance After Polypectomy: A Consensus Update by the US Multi-Society Task Force on Colorectal Cancer and the American Cancer Society. *Gastroenterology* **130**, 1872–1885 (2006).

164. Gupta, S. *et al.* Recommendations for Follow-Up After Colonoscopy and Polypectomy: A Consensus Update by the US Multi-Society Task Force on Colorectal Cancer. *Gastrointest. Endosc.* **91**, 463-485.e5 (2020).

165. Lieberman, D. A. *et al.* Guidelines for Colonoscopy Surveillance After Screening and Polypectomy: A Consensus Update by the US Multi-Society Task Force on Colorectal Cancer. *Gastroenterology* **143**, 844–857 (2012).

166. Goede, S. L. *et al.* Cost-savings to Medicare from Pre-Medicare Colorectal Cancer Screening. *Med. Care* **53**, 630–638 (2015).

167. van Hees, F. *et al.* The Appropriateness of More Intensive Colonoscopy Screening than Recommended in Medicare Beneficiaries: A Modeling Study. *JAMA Intern. Med.* **174**, 1568–1576 (2014).

168. Ploskas, N., Laughman, C., Raghunathan, A. U. & Sahinidis, N. V. Optimization of circuitry arrangements for heat exchangers using derivative-free optimization. *Chem. Eng. Res. Des.* **131**, 16–28 (2018).

169. Ackley, D. *A Connectionist Machine for Genetic Hillclimbing*. (Springer US, 2012).

170. Young, D., Haney, W. & Cremaschi, S. Derivative-free optimization of combinatorial problems – A case study in colorectal cancer screening. *Comput. Chem. Eng.* **145**, 107193 (2021).

171. Neumann, P. J., Cohen, J. T. & Weinstein, M. C. Updating cost-effectiveness - The curious resilience of the $50,000-per-QALY threshold. *N. Engl. J. Med.* **371**, 796–797 (2014).

172. Meester, R. G. S. *et al.* Optimizing colorectal cancer screening by race and sex: Microsimulation analysis II to inform the American Cancer Society colorectal cancer screening guideline. *Cancer* **124**, 2974–2985 (2018).

173. Center for Disease Control and Prevention. Colorectal Cancer Screening Tests. (2020). Available at: https://www.cdc.gov/cancer/colorectal/basic_info/screening/tests.htm. (Accessed: 10th February 2021)

174. Glover, F. & Woolsey, E. Further Reduction of Zero-One Polynomial Programming Problems to Zero-One Linear Programming Problems. *Oper. Res.* **21**, 156–161 (1973).

175. Oral, M. & Kettani, O. A linearization procedure for quadratic and cubic mixed-integer problems. *Oper. Res.* **40**, S109--S116 (1992).

176. Arias, E. United States life tables, 2000. *Natl. Vital Stat. Reports* **51**, 1–63 (2002).

177. Young, D. & Cremaschi, S. A Simulation-based Optimization Approach to Develop Personalized Colorectal Cancer Screening Strategies. in *Computer Aided Chemical Engineering* **44**, 2125–2130 (Elsevier, 2018).

178. American Cancer Society. Colorectal cancer Facts & Figures 2020-2022. (2020).

179. Kahn, H. & Marshall, A. W. Methods of reducing sample size in Monte Carlo computations. *J. Oper. Res. Soc. Am.* **1**, 263–278 (1953).

180. Brooks, S. Markov chain Monte Carlo method and its application. *J. R. Stat. Soc. Ser. D (the Stat.* **47**, 69–100 (1998).

181. Owen, A. & Zhou, Y. Safe and effective importance sampling. *J. Am. Stat. Assoc.* **95**,

135–143 (2000).

182. Papavasiliou, A. & Oren, S. S. Multiarea stochastic unit commitment for high wind penetration in a transmission constrained network. *Oper. Res.* **61**, 578–592 (2013).

183. Morokoff, W. J. & Caflisch, R. E. Quasi-Monte Carlo Integration. *J. Comput. Phys.* **122**, 218–230 (1995).

184. Joseph, V. R. Space-filling designs for computer experiments: A review. *Qual. Eng.* **28**, 28–35 (2016).

185. Yadav, J. & Sharma, M. A Review of K-mean Algorithm. *Int. J. Eng. trends Technol.* **4**, 2972–2976 (2013).

186. Yuan, C. & Yang, H. Research on K-value selection method of K-means clustering algorithm. *J* **2**, 226–235 (2019).

187. Satopää, V., Albrecht, J., Irwin, D. & Raghavan, B. Finding a 'kneedle' in a haystack: Detecting knee points in system behavior. *Proc. - Int. Conf. Distrib. Comput. Syst.* 166–171 (2011). doi:10.1109/ICDCSW.2011.20

188. Pelleg, D., Moore, A. W. & others. X-means: Extending k-means with efficient estimation of the number of clusters. in *Icml* **1**, 727–734 (2000).

189. Audet, C., Denni, J., Moore, D., Booker, A. & Frank, P. A surrogate-model-based method for constrained optimization. in *8th symposium on multidisciplinary analysis and optimization* 4891 (2000).

190. Han, Z.-H., Zhang, K.-S. & others. Surrogate-based optimization. *Real-world Appl. Genet. algorithms* **343**, (2012).

191. Haykin, S. *Neural networks and learning machines, 3/E*. (Pearson Education India, 2009).

192. Eckle, K. & Schmidt-Hieber, J. A comparison of deep networks with ReLU activation

function and linear spline-type methods. *Neural Networks* **110**, 232–242 (2019).

193. Kunanbayev, K., Temirbek, I. & Zollanvari, A. Complex Encoding. in *2021 International Joint Conference on Neural Networks (IJCNN)* 1–6 (2021).

194. Anderson, R., Huchette, J., Ma, W., Tjandraatmadja, C. & Vielma, J. P. Strong mixed-integer programming formulations for trained neural networks. *Math. Program.* **183**, 3–39 (2020).

195. Breiman, L. Random forests. *Mach. Learn.* **45**, 5–32 (2001).

196. Burman, P. A comparative study of ordinary cross-validation, v-fold cross-validation and the repeated learning-testing methods. *Biometrika* **76**, 503–514 (1989).

197. Heaton, J. *Introduction to neural networks with Java*. (Heaton Research, Inc., 2008).

Appendix A Adenoma transition probability by size, age, sex, and location

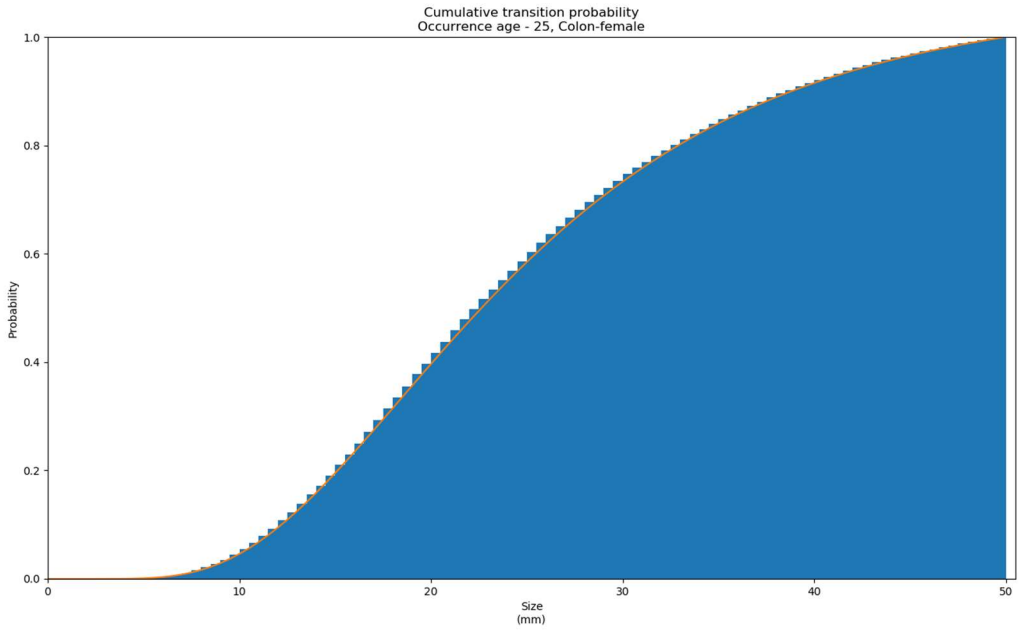

Figure A.5.1 Transition probability of an adenoma by size initiating at age 25 in the colon of a female

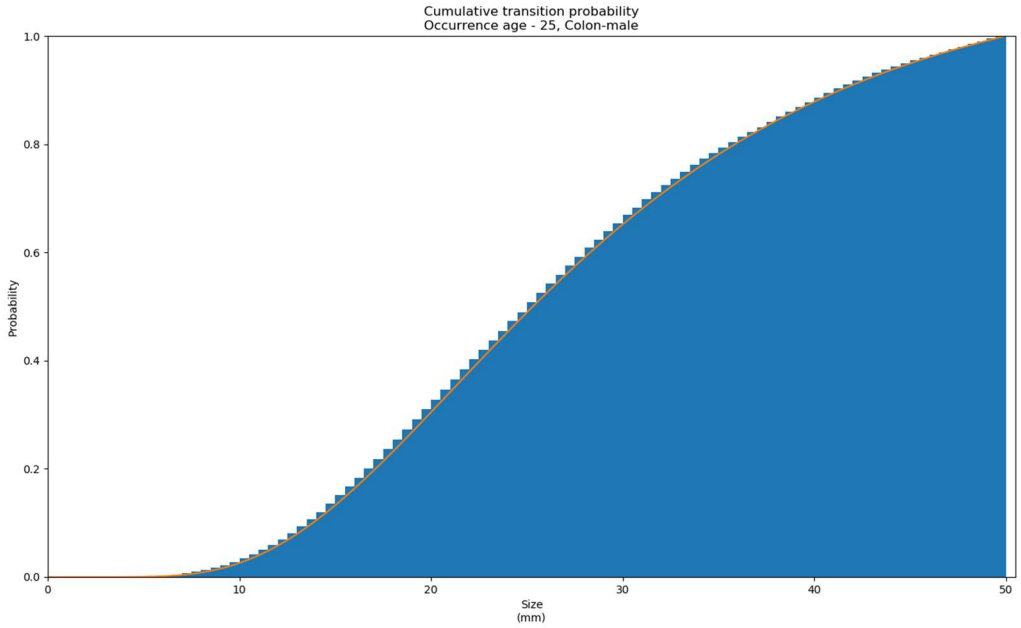

Figure A.5.2 Transition probability of an adenoma by size initiating at age 25 in the colon of a male

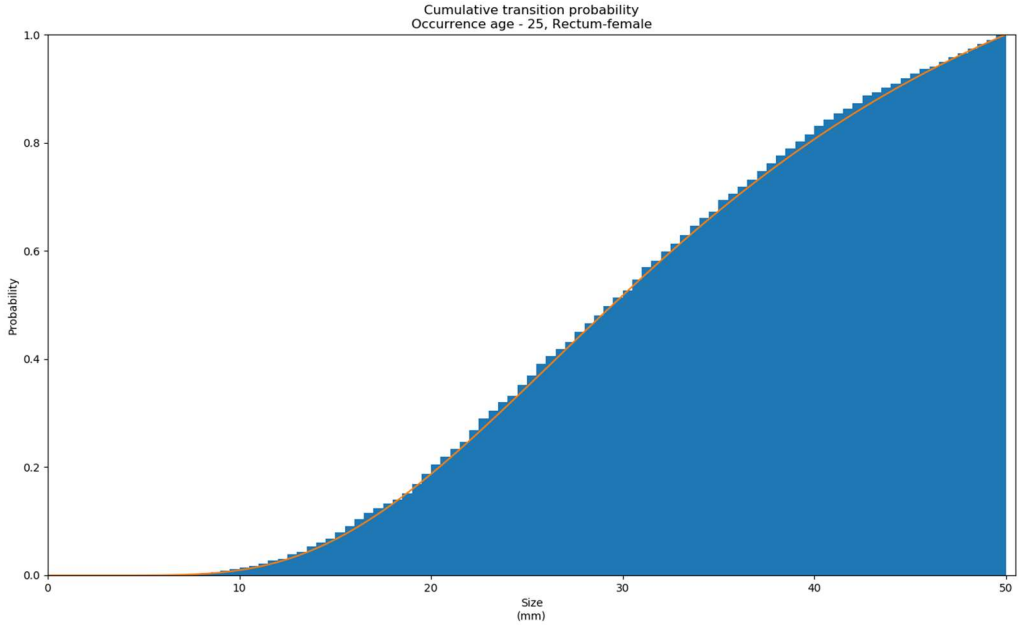

Figure A. 5.3 Transition probability of an adenoma by size initiating at age 25 in the rectum of a female

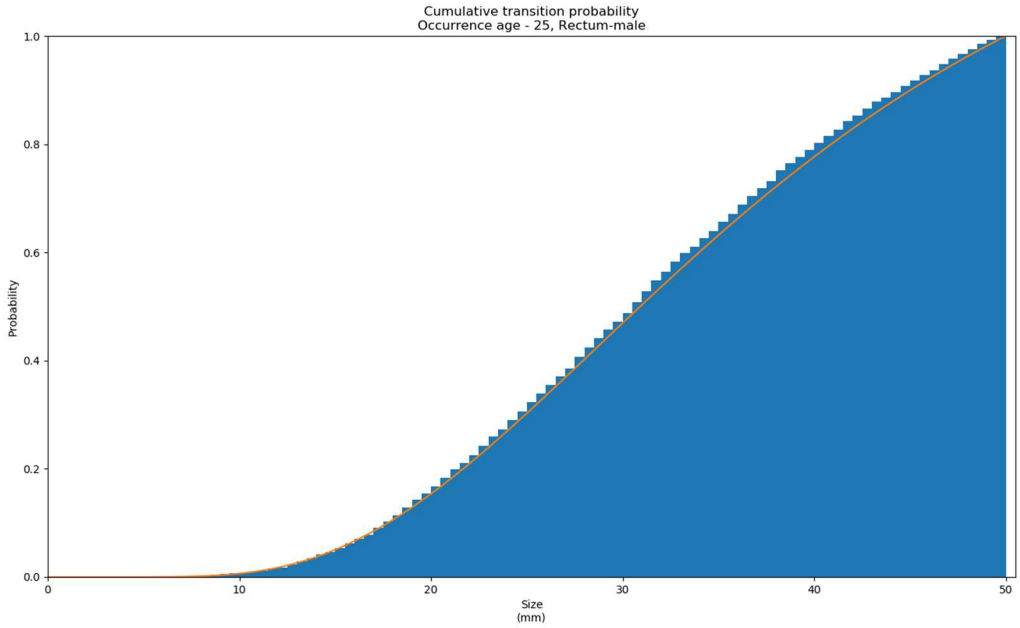

Figure A.5.4 Transition probability of an adenoma by size initiating at age 25 in the rectum of a male

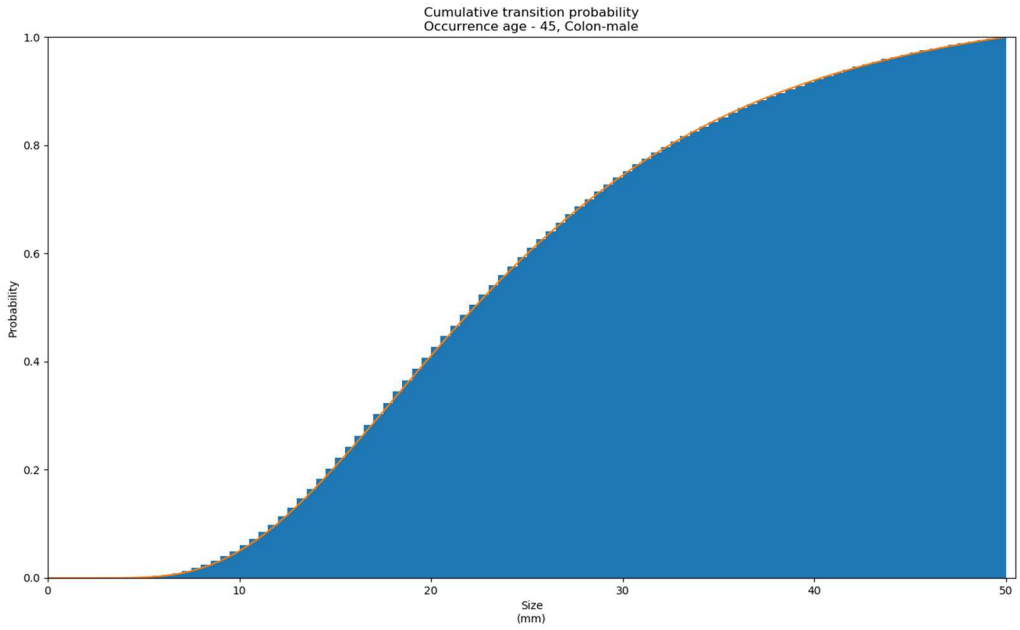

Figure A.5.5 Transition probability of an adenoma by size initiating at age 45 in the colon of a male

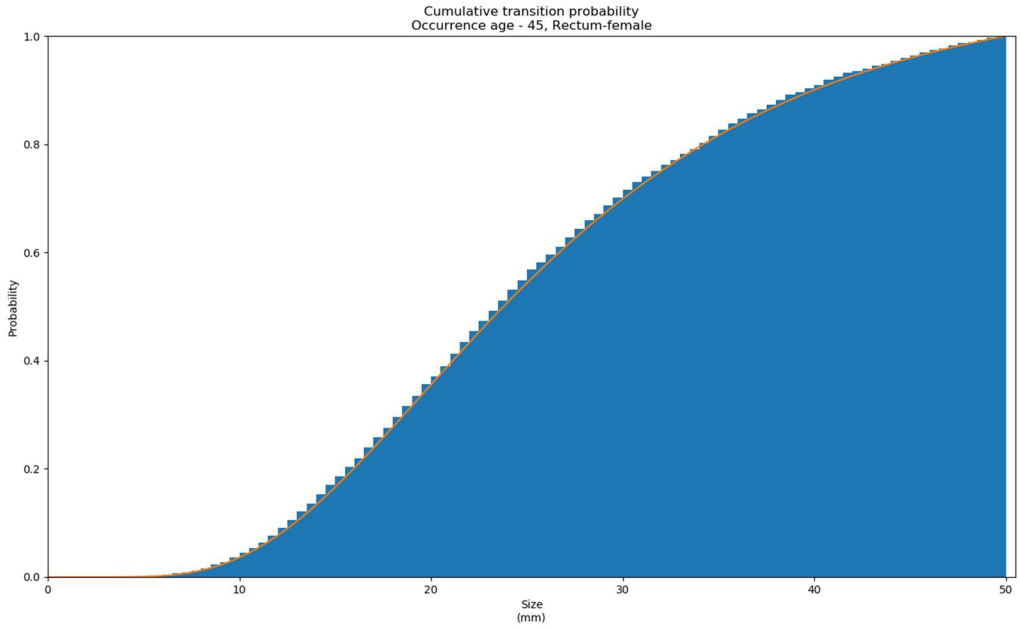

Figure A.5.6 Transition probability of an adenoma by size initiating at age 45 in the rectum of a female

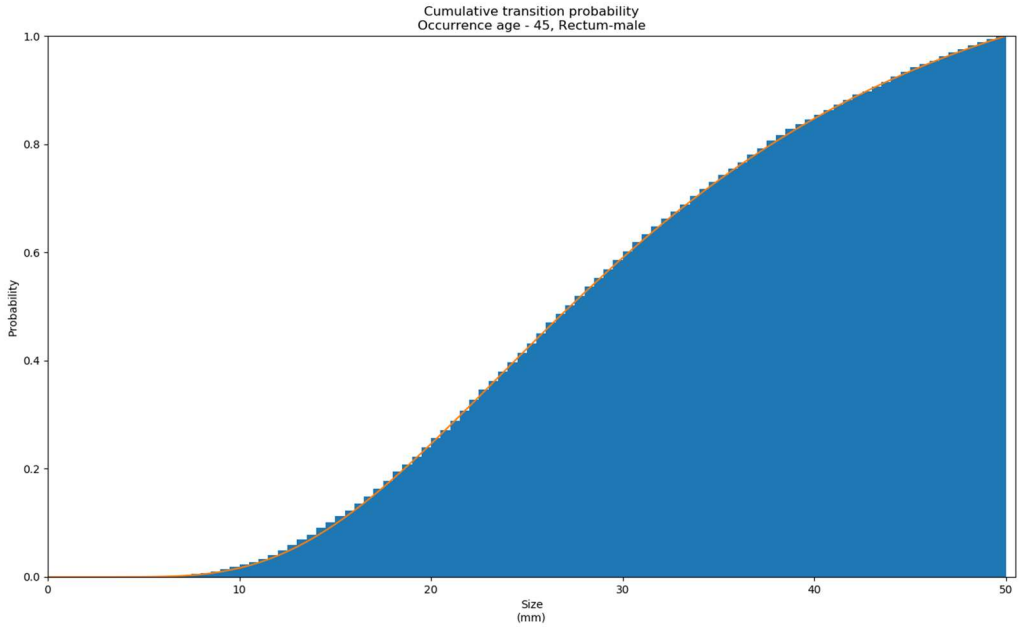

Figure A5.7 Transition probability of an adenoma by size initiating at age 45 in the rectum of a male

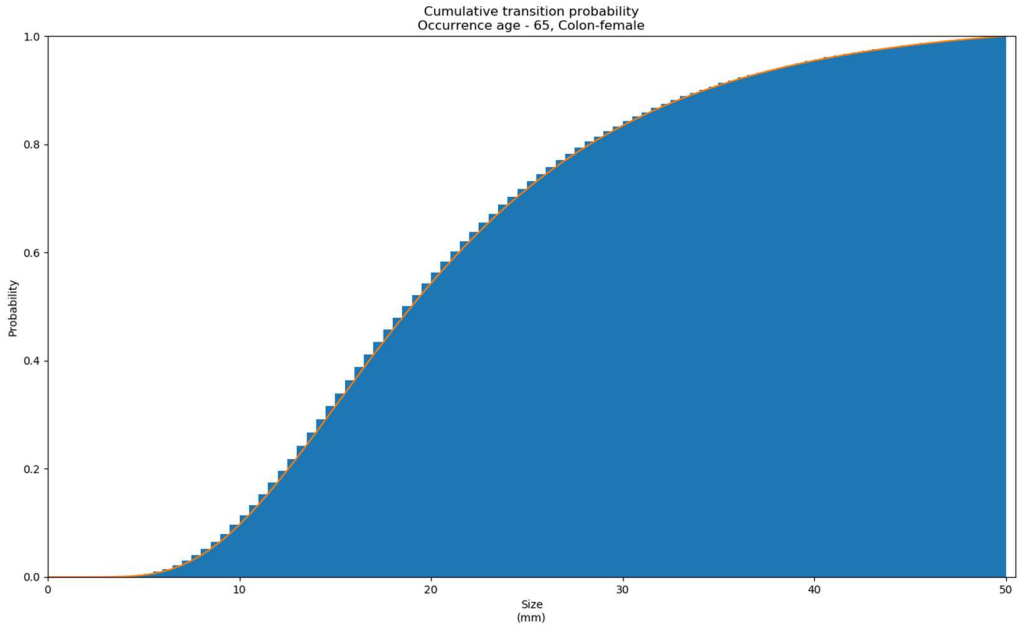

Figure A.5.8 Transition probability of an adenoma by size initiating at age 65 in the colon of a female

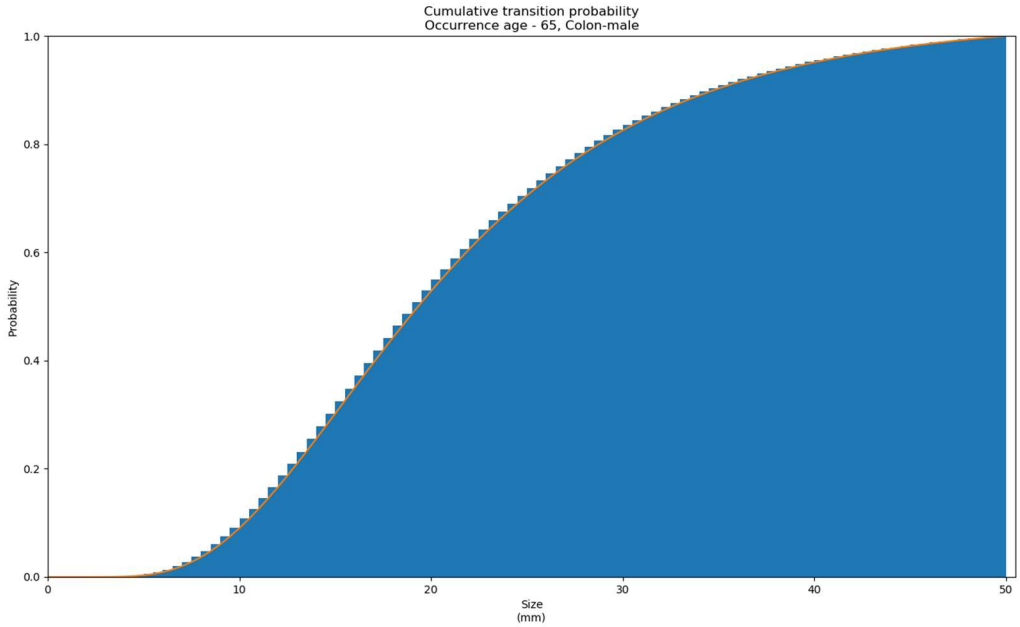

Figure A.5.9 Transition probability of an adenoma by size initiating at age 65 in the colon of a male

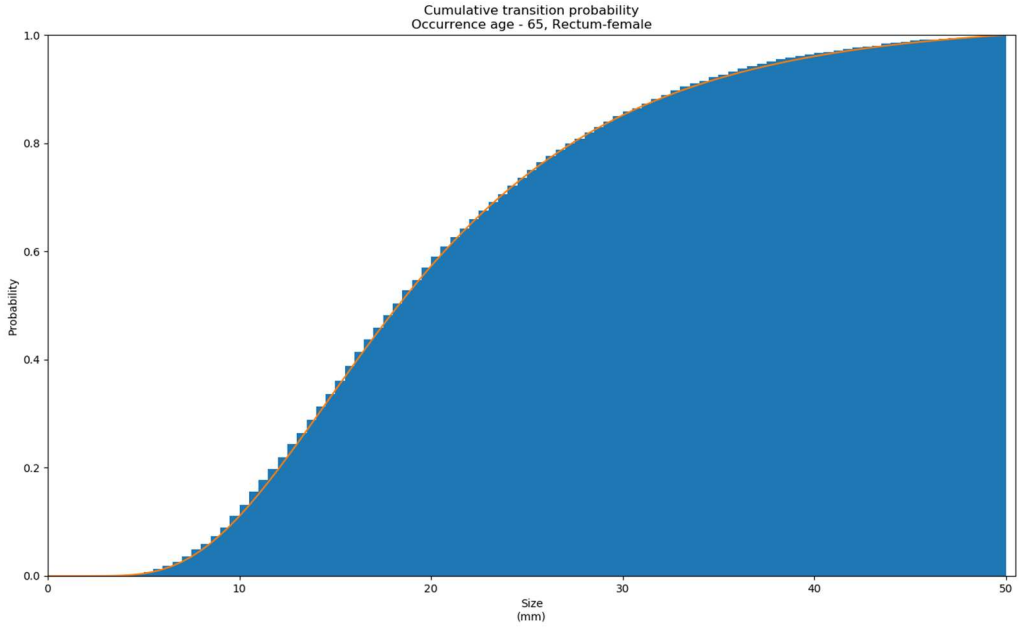

Figure A.5.10 Transition probability of an adenoma by size initiating at age 65 in the rectum of a female

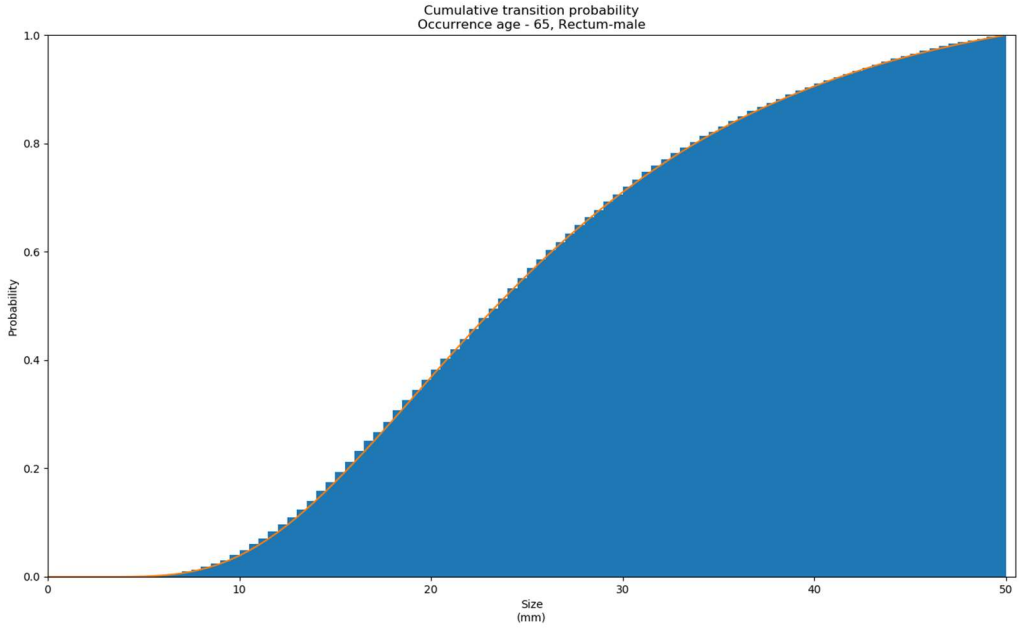

Figure A.5.11 Transition probability of an adenoma by size initiating at age 65 in the rectum of a male

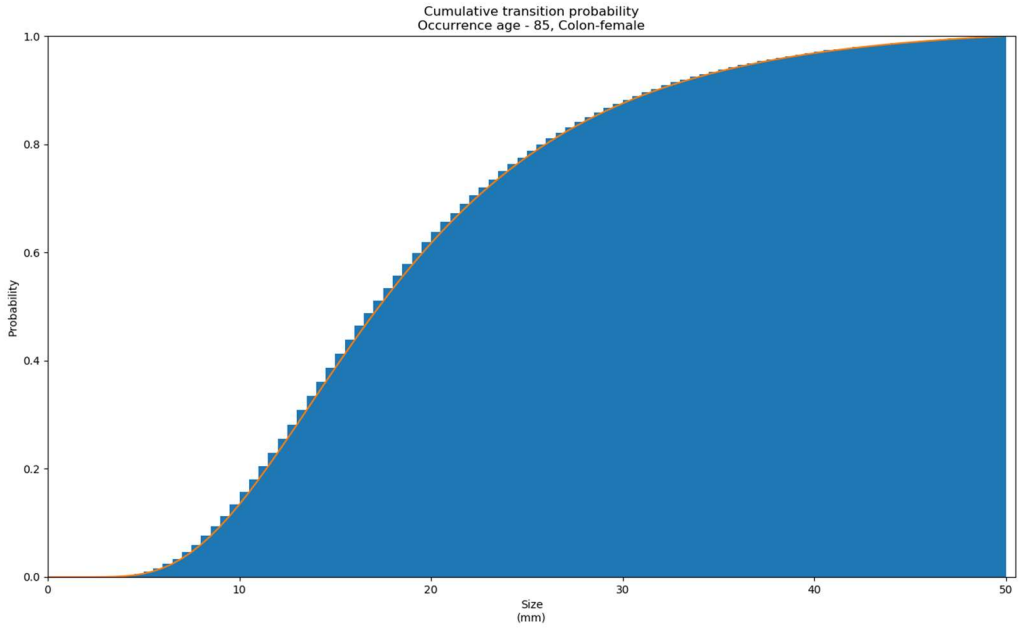

Figure A.5.12 Transition probability of an adenoma by size initiating at age 85 in the colon of a female

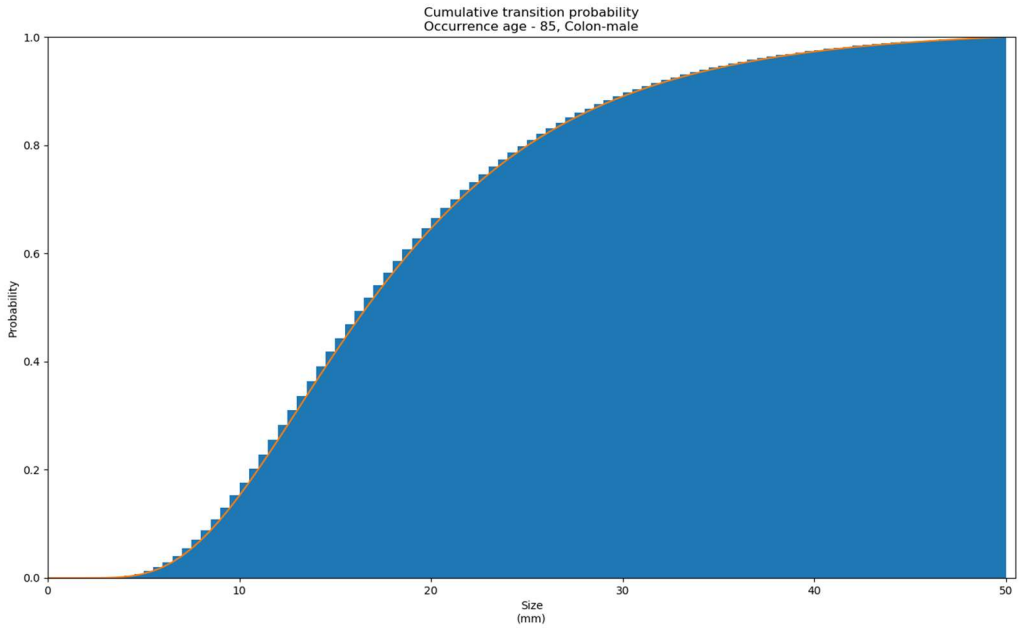

Figure A.5.13 Transition probability of an adenoma by size initiating at age 85 in the colon of a male

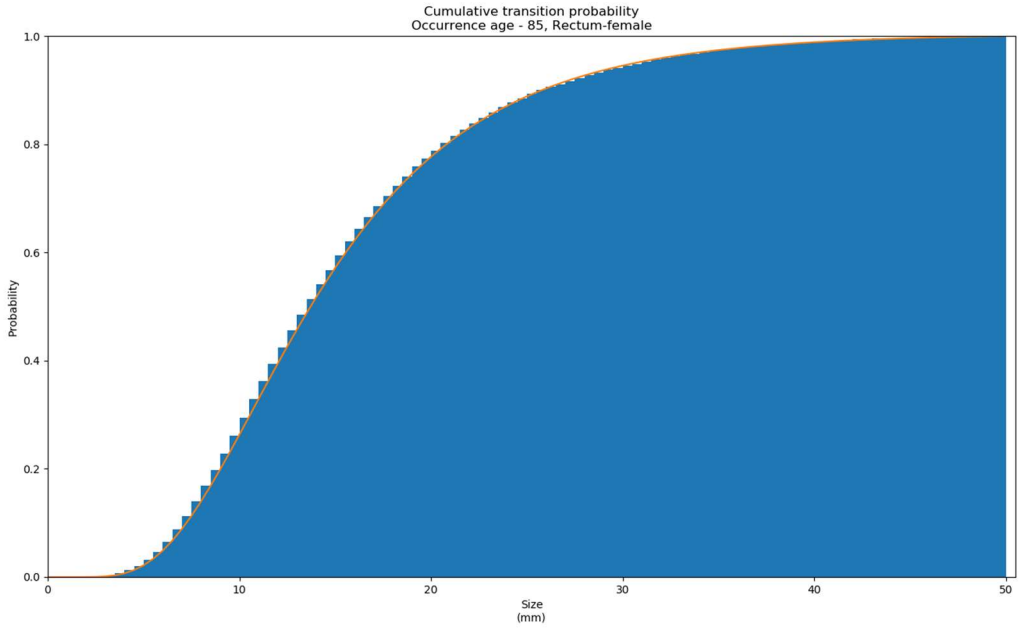

Figure A.5.14 Transition probability of an adenoma by size initiating at age 85 in the rectum of a female

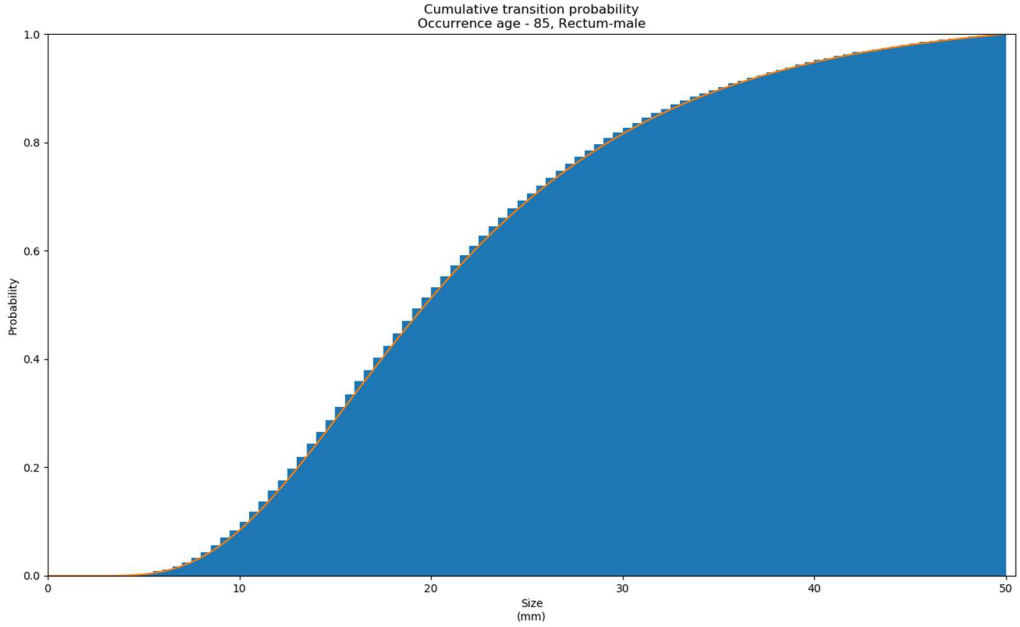

Figure A.5.15 Transition probability of an adenoma by size initiating at age 85 in the rectum of a male

Appendix B Machine learning model parameter determination

The parameters of both the NN and RF were determined using 5-fold cross-validation[196] with 15 Monte Carlo replications, withholding 20 % of the training data to be used as testing data. The metric used to evaluate the model quality is the average root mean square error (RMSE) of the testing data, to avoid overfitting.

B.1 Neural Network parameter determination

With ANN, the main parameters to be determined are 1) the number of hidden layers, 2) the number of neurons per layer, 3) the number of training epochs for determining the weights and biases of the links in the network, and 4) the activation function. Since the number of layers and the activation functions were automatically chosen to reduce the complexity of the model, parameter determination for the ANN only required the evaluation of the number of neurons and the number of training epochs.

First the number of neurons was determined. There have been a number of rule of approaches for determining the number of neurons within a hidden layer, including twice the size of the input layer, somewhere between the size of the input and size of the output layer, and the sum of the size of the input and size of the output layers[197]. The size of the input layer for the neural networks is 24, where two variables represent the screening ages, and 22 represent the one-hot encoding of the scenario realizations. The output layer of the neural network has only one variable, the screening benefit metric, either QALY lived or QALY gained. Combining the rules of thumb led to a parameter search space of $[2, 50]$ for the number of neurons in the hidden layer. The number of training epochs was set to be an arbitrarily large value of 1,500, and was terminated early if no improvement of the testing RMSE value is seen for 10 consecutive epochs. Multiple

NNs are trained increasing the number of neurons in the hidden layer from 2 to 50 in increments of two. The resulting average RMSE for the trained NNs are presented in Figure B.1, where plot A shows the parameter evaluation for QALY lived and plot B shows the parameter evaluation for QALY gained for both the training and testing data. From the plot, the value of the parameter is chosen to be the value at which the decrease in RMSE begins to level out. The number of neurons was determined to be from 34 to 38 and 20 to 24 for QALY lived and QALY gained respectively.

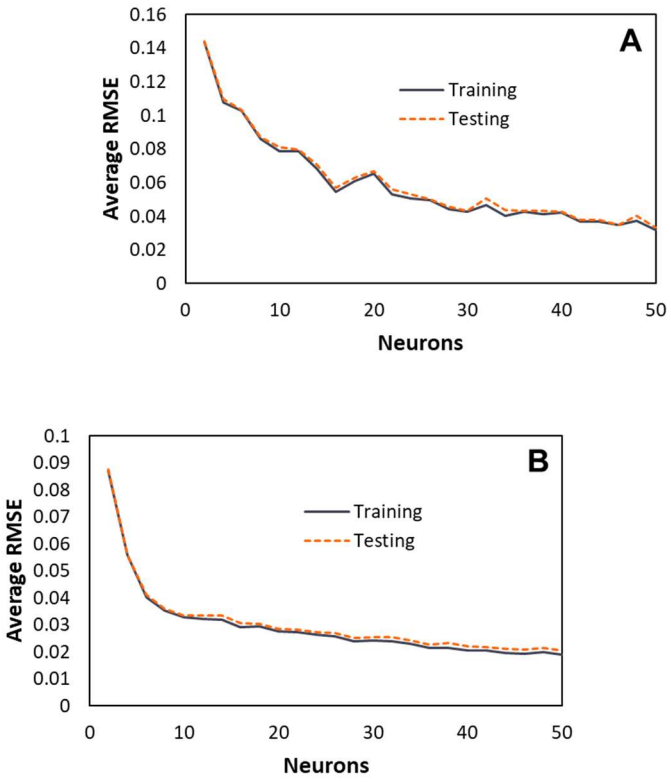

Figure B.1 Number of neuron parameter evaluation showing the model's RMSE as a function of the number of neurons in the hidden layer for the screening benefit metric of A) QALY lived and B) QALY gained.

Next the number of training epochs was needed to be determined. A number of NNs were then trained incrementing the number of neurons by one from 34 to 38 and 20 to 24 for QALY lived and QALY gained respectively. The testing RMSE is recorded as a function of the number of training epochs from 100 to 950 in increments of 50 epochs. The resulting average RMSE for

the trained NNs are presented in Figure B.2, where plot A shows the parameter evaluation for QALY lived and plot B shows the parameter evaluation for QALY gained. Using the same approach for determining the approximate number of neurons, the number of training epochs and neurons was determined visually from the plot leading to 34 neurons and 800 training epochs for QALY lived and 23 neurons and 700 training epochs for QALY gained.

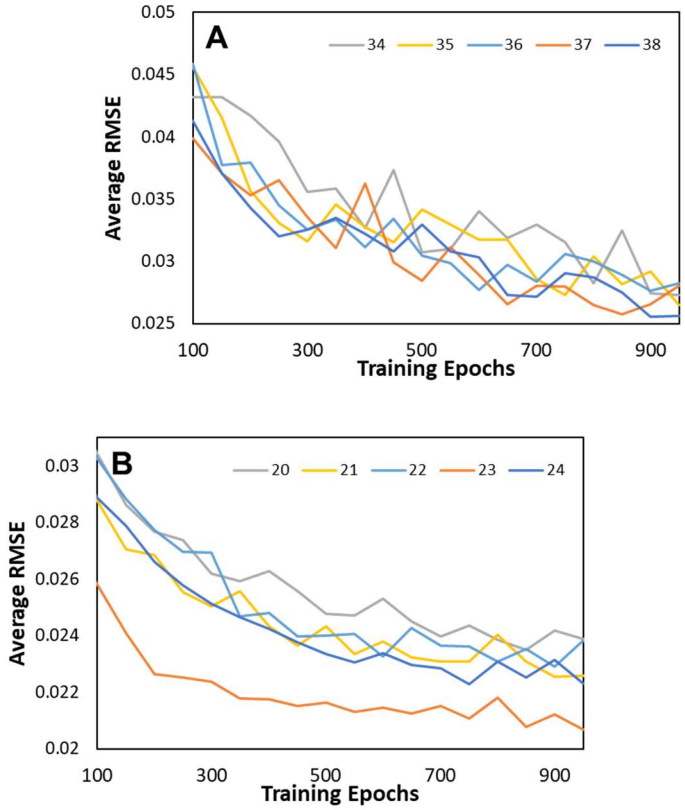

Figure B.2 Training epoch parameter evaluation showing the model's RMSE as a function of the training epochs for the screening benefit metric of A) QALY lived and B) QALY gained.

B.2 Random Forest parameter determination

The parameter to be determined for the RF model are 1) the number of trees within the RF and 2) the maximum depth, or number of branching points, of the RF. As with the NN models, the parameter values were determined sequentially, first with the number of trees followed by the maximum tree depth. Unlike with NN, we evaluate the incorporation of the scenarios into the

model by comparing the use of a total of 22 different RF models, one for each scenario, (RF22 method) versus a single RF model using one-hot encoding (RF1H method) to include the scenarios as input variables.

First, both the number of trees to include in the model and if the scenarios should be represented by a numerous individual RF models or a single large RF model were evaluated together. To evaluate the number trees to use in the RF, a number of RF models were trained starting with five trees and increasing the number of trees by five up to a total of 225 trees. This was done for both the RF22 method and RF1H method. To compare the two methods, the average testing RMSE of the 22 models in the RF22 method is compared to the average testing RMSE of the RF1H method. The resulting testing RMSE as a function of the number of trees in the RF is presented in Figure B.3. From the results, it is clear that the RF22 method provides a more accurate set of regression models than constructing a single model. Additionally from the plot, it can be seen that for both cases there is no clear improvement in the models after approximately 75 and 75 trees for QALY lived and QALY gained respectively.

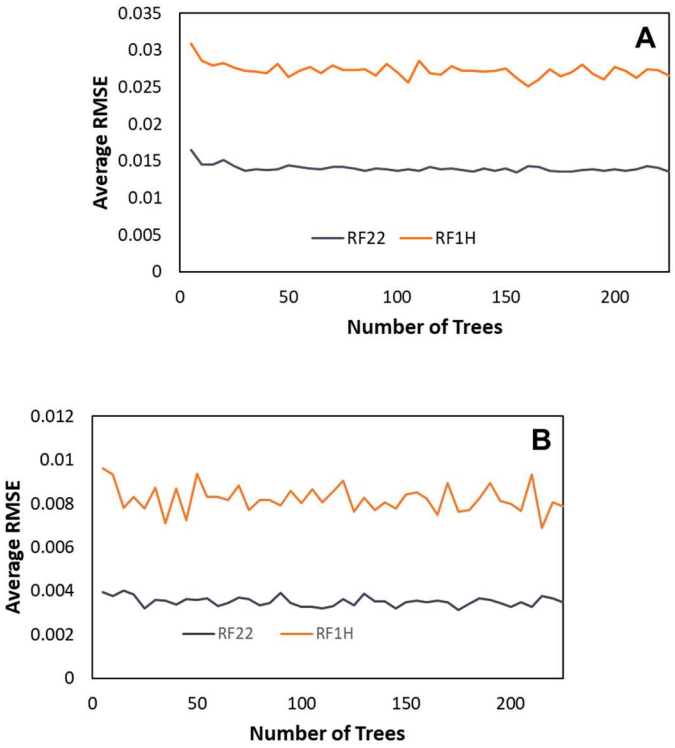

Figure B.3 Number of trees parameter evaluation showing the model's RMSE as a function of the number of trees in the forest model for the screening benefit metric of A) QALY lived and B) QALY gained.

The number of trees in the forest was then set based visually from Figure B.3 to assess the max depth of each tree in the RF model. The maximum depth allowed for each tree was increased from one to 50, the resulting testing RMSE is presented in Figure B.4. Visually it was determined that the max depth needed for the model was a depth of 9 and 10 for QALY lived and QALY gained respectively.

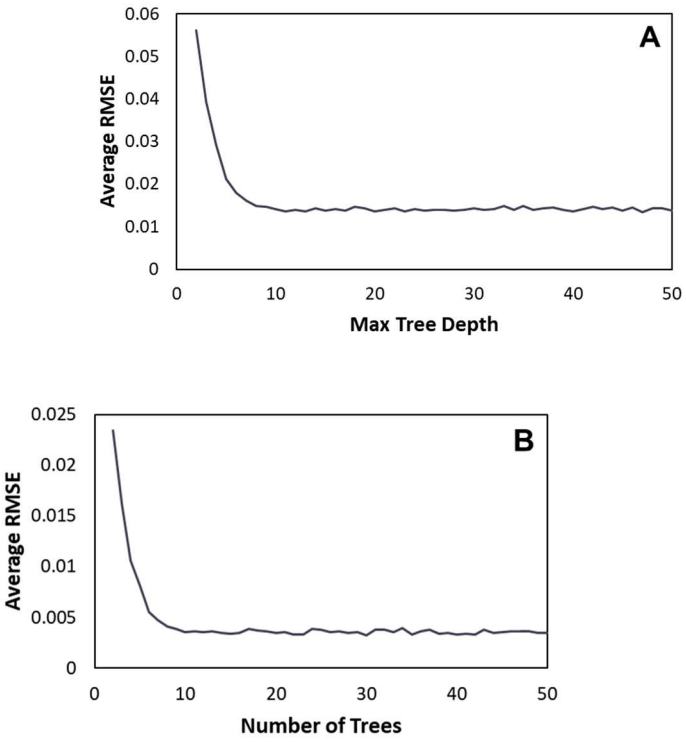

Figure B.4 Maximum tree depth parameter evaluation showing the model's RMSE as a function of the max depth of each tree in the model for the screening benefit metric of A) QALY lived and B) QALY gained.