Optimising the use of genetic testing in prevention of CHD using Decision Programming

Helmi Hankimaa

School of Science

Bachelor's thesis
Espoo 17.03.2021

Supervisor

Prof. Fabricio Oliveira

Advisor

M.Sc.(Tech.) Olli Herrala

Aalto University
School of Science

Aalto University
School of Science

Aalto University, P.O. BOX 11000, 00076 AALTO
www.aalto.fi
Abstract of the bachelor's thesis

Author Helmi Hankimaa

Title Optimising the use of genetic testing in prevention of CHD using Decision Programming

Degree programme Engineering Physics and Mathematics

Major Mathematics and Systems Sciences **Code of major** SCI3029

Teacher in charge Prof. Fabricio Oliveira

Advisor M.Sc.(Tech.) Olli Herrala

Date 17.03.2021 **Number of pages** 25 **Language** English

Abstract

The current system of providing preventative care for coronary heart disease (CHD) can be seen as unsuccessful as the majority of CHD events occur in the population that is not receiving preventative treatment. Currently, treatment is administered based on the patient's estimated risk of suffering a CHD event. The genetic risk score has been developed to improve CHD risk stratification. Since the genetic risk score is more expensive than traditional testing methods, cost-benefit analysis is needed to evaluate the value of it in preventative care decision making of CHD.

The aim of this study is to use the Decision Programming framework to develop an optimal decision strategy for allocating preventative care for CHD. This study develops an optimal testing and treatment strategy with the objective of maximising net monetary benefit. The problem is solved using the Decision Programming framework which combines aspects of stochastic programming and decision analysis. The optimisation problem formulation relies on an influence diagram representation of the problem. The initial formulation was computationally intractable. In order to obtain a solution, the problem was decomposed into a set of 101 subproblems.

By analysing and combining the optimal solutions of the subproblems an optimal decision strategy was successfully developed. The optimal strategy is a two-stage screening process of traditional risk score and genetic risk score testing. The Decision Programming framework was applicable to the problem because it allowed modeling endogenous uncertainties, a multi-staged decision strategy and the multiple value nodes in the influence diagram.

Keywords Decision programming, endogenous uncertainty, influence diagrams, decision analysis, stochastic programming, mixed integer linear programming, preventative care allocation

Contents

1 Introduction

Coronary heart disease (CHD) is a type of cardiovascular disease. Its prevalence has decreased in Finland since the 1960's due to improvements in lifestyle, preventative care and treatment methods. Despite this, according to THL (2020), cardiovascular diseases were still the number one cause of death in Finland in 2018.

Improving the system of providing preventative care for CHD is key in reducing its prevalence. Statin treatment is a form of preventative care, which reduces levels of cholesterol in blood vessels. This treatment is provided to patients with a high (20%) 10-year absolute risk of cardiovascular disease. The risk estimate for a patient is based on the traditional risk score (TRS), which relies heavily on the main traditional risk factors. However, despite the traditional risk factors being very prevalent in the population who develop CHD, Weissler (2004) showed that the predictive power of these risk factors is poor. According to Tikkanen et al. (2013), more than half of CHD events occur within the population that has not been classified to be at high risk and thus, are not receiving preventative care.

In response, new tests for predicting a patient's risk of CHD have been developed. The genetic risk score (GRS) is one of them. These tests provide more information and may enable health care providers to detect more patients at high risk and provide preventative care for them. Thus, more lives are saved. However, all of this incurs additional costs and an increased workload for the health care system. Thus, it is essential to develop a decision strategy that yields the most benefit. A decision strategy in this context is a plan that dictates which patients are tested and treated and which tests are used. Cost-benefit analysis is used to analyse and develop decision strategies.

The aim of this thesis is to evaluate if the Decision Programming framework is suitable for the cost-benefit analysis performed by Hynninen et al. (2019). In their study, Hynninen et al. (2019) solved for an optimal testing and statin treatment allocation strategy for CHD using dynamic programming. In this paper, the same problem will be modeled using the Decision Programming framework, recently developed by Salo et al. (2019). The strengths of the framework and challenges presented by using it will be discussed. In the model, the available diagnostic tests are TRS and GRS and the cost-benefit objective is evaluated as a weighted sum of testing costs and health outcomes.

The structure of this thesis is as follows. Prior research on developing decision strategies in health care and cost-benefit analyses of these strategies is reviewed in Section 2. The thesis continues to describe the methodology of Decision Programming and the development of the optimisation model in Section 3. The results of the optimisation are presented in Section 4. Further discussion of the results and the conclusions that can be drawn from them are described in Section 5.

2 Literature review

The cost-benefit study by Hynninen et al. (2019) solved for the optimal testing and statin allocation strategy from the health care system's perspective. The model determined which patients should be tested and treated based on their risk of CHD. Furthermore, it determined the testing strategy that should be used in the risk stratification. Thus, their model developed a decision strategy and simultaneously evaluated it in terms of cost-benefit. Previously, these two aspects of preventative care decision making have not been researched in a study simultaneously. The general guide of how cost-benefit analyses in this area have been conducted is that a set of predefined decision strategies is chosen, these strategies are simulated and the results are compared. The strategies are generally adopted from epidemiological studies which have researched the effectiveness of the testing methods in risk stratification or from national health care guidelines.

Therefore, the first section of this literature review focuses on how testing and treatment strategies have previously been developed. The second section provides a review of studies that have analysed cost-benefit or cost-effectiveness of decision strategies in the area of allocating preventative care for cardiovascular diseases.

2.1 Testing and treatment strategy development

One of the first frameworks for medical decision making of whether to treat, test or not treat was developed by Pauker and Kassirer (1980). This framework provides an analytical basis for testing and treatment strategies. They developed two thresholds, the 'testing' threshold and the 'test-treatment' threshold. The thresholds are probability cut-offs and they divide subjects into three groups: if the risk of disease is below the 'testing' threshold treatment should be withheld, if it is above the 'test-treatment' threshold treatment should be given and if the risk falls in between these thresholds then the diagnostic test should be performed and the treatment decision made based on its results. The thresholds are visualised in Figure 1. The framework of Pauker and Kassirer (1980) only considers the patient's welfare and does not consider economic aspects.

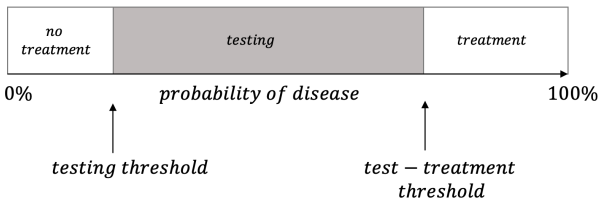

Figure 1: Visualisation of the 'testing' and 'test-treatment' thresholds. Figure adapted from Pauker and Kassirer (1980).

An example of a medical study where these threshold values were defined for CHD is the study on the effectiveness of GRS by Tikkanen et al. (2013). They

tested a hypothesis that a two stage screening process of TRS and then GRS would improve risk classification for CHD. They found that after performing TRS on all subjects, performing GRS on the intermediate risk scores (10%-20%) produced a significant improvement in the classification. They saw that 31% in the intermediate risk group were correctly reclassified into low- and high-risk categories as a result of the GRS testing. This improvement in classification enables better preventative treatment allocation and Tikkanen et al. (2013) estimated that it would prevent one additional CHD event over 14 years for every 135 people screened. This estimate was based on the fact that according to current guidelines, patients with a risk greater than 20% receive preventative statin treatment. Based on their study, the 'testing' and 'test-treatment' thresholds for the GRS are 10% and 20%. These values were determined based on the risk levels for which the GRS provides the most information gain and the current health care guidelines. The testing strategy developed by Tikkanen et al. (2013) was used as a reference strategy in the study by Hynninen et al. (2019).

Hynninen et al. (2019) developed a testing and treatment strategy by optimising a cost-benefit objective consisting of the health outcomes and testing costs. This included them determining the 'testing' and 'test-treatment' thresholds for TRS and GRS from the perspective of net monetary benefit (NMB). The threshold values for GRS were different than the ones found in the study by Tikkanen et al. (2013). This is due to the different perspectives – pure patient welfare versus NMB – that the studies were conducted from. For example, the national health care guidelines for allocating treatment were not considered in the optimisation by Hynninen et al. (2019). This showcases that the thresholds described by Pauker and Kassirer (1980) are not absolute for a test for a given disease because the perspective of the study affects the found threshold values.

2.2 Cost-benefit analysis of decision strategies in the context of cardiovascular diseases

New testing and treatment strategies are evaluated from an economical perspective in cost-benefit analyses. Cost-benefit analyses often examine a set of decision strategies, which they compare to find the optimal strategy. These studies often use decision trees to express the strategies and Markov state-transition models to simulate them. This approach was adopted by Roberts et al. (2015), Wordsworth et al. (2010) and Jarmul et al. (2018). The simulation parameters are set based on data from health records. The results of the simulation are used to compare the strategies.

Roberts et al. (2015) used a decision tree and Markov model methodology to study the costs and benefits of using coronary artery calcium (CAC) testing for CHD risk stratification for intermediate risk patients. In their study, they defined four strategies: treat all, treat according to current guidelines, treat if CAC ≥ 1, and treat if CAC ≥ 100. They found that if the benefit was measured as CHD events averted, then the strategy of treating everyone with CAC ≥ 1 is the most cost-effective. However, if the benefit was measured in quality-adjusted life-years (QALYs) then the strategy where only patients with CAC ≥ 100 is favoured. The

reason behind this was that QALYs accounted for patients' disfavor and possible disutility due to side-effects of the statin treatment. This showcases the importance of choosing the measure of benefit intelligently in cost-effectiveness and cost-benefit analysis.

Wordsworth et al. (2010) performed another cost-effectiveness study using a decision tree and a Markov model. They compared the cost-effectiveness of genetic and clinical screening strategies of family members in risk stratification of hypertrophic cardiomyopathy, which is another type of cardiovascular disease. The genetic testing strategy was found to be more costly due to more high-risk individuals being detected and thus accumulated more preventative treatment costs. However, with genetic testing the estimated life years saved was higher. The genetic testing strategy was determined to be more cost-effective, meaning that the yielded benefits outweighed incurred costs. This was determined from the fact that the incremental cost per life year saved was below the willingness-to-pay threshold. The willingness-to-pay threshold describes how much society is willing to pay for a life year saved. Whether or not the the life years have been adjusted with the quality of life depends on the study. The result of Wordsworth et al. (2010) exemplifies the need for cost-effectiveness analysis, because comparing the costs and their effects is not always straightforward. The more expensive strategy might be the more cost-effective one.

Jarmul et al. (2018) performed cost-effectiveness analysis in a similar context as Hynninen et al. (2019). Jarmul et al. (2018) analysed the value of using a cardiovascular genetic risk score (cGRS) in targeting statin therapy for primary prevention of atherosclerotic cardiovascular diseases (ASCVD). Coronary heart disease is a type of ASCVD. The study focused on individuals at low-to-intermediate (2,5%-7,5%) risk because based on expert opinion, preventative care in this context should be allocated to those with a risk $\geq$ 7,5%. Their model was also a state-transition Markov model. By changing parameters describing patient profiles and testing and treatment plans, they were able to compare ASCVD incidence, quality of life, mortality and costs for the different strategies. The primary measure of benefit was gained QALYs. They defined four testing and treatment strategies: (1) test none and treat all, (2) test none and treat none, (3) test all and treat if cGRS is high, and (4) test all and treat if cGRS is intermediate or high. The authors found that using cGRS was not cost-effective in allocating statin treatment for people with low-to-intermediate risk. The shortcoming of this study is that only individuals at low-to-intermediate risk were studied. Furthermore, the cGRS 'testing' and 'test-treatment' thresholds were predefined as 2,5% and 7,5%. It was not explored, whether 2,5% and 7,5% are the best 'testing' and 'test-treatment' thresholds for cGRS. The study by Hynninen et al. (2019) improved on these aspects.

This paper replicates the study by Hynninen et al. (2019). In their study, six testing and treatment strategies were evaluated. However, in contrast to the studies outlined above, in these strategies the allocation of tests and treatment according to risk levels was not predefined. The optimal allocation was instead obtained from the dynamic programming model used in the study. The six strategies were

1. no tests and no treatment ('No treatment'),

2. using prior risk to allocate treatment ('Treatment optimised'),
3. performing TRS on optimised patient segment and allocating treatment based on updated risk estimates ('TRS optimized'),
4. performing GRS on optimised patient segment and allocating treatment based on updated risk estimates ('GRS optimized'),
5. performing TRS on optimised patient segment and based on its results performing GRS optimally to allocate treatment ('TRS & GRS optimized'),
6. performing GRS on optimised patient segment and based on its results performing TRS optimally to allocate treatment ('GRS & TRS optimized').

They also used four reference strategies:

7. performing TRS on all patients ('TRS for all'),
8. performing GRS on all patients ('GRS for all'),
9. performing TRS for all patients and then GRS for the ones with an updated risk estimate between 10%-20% (TRS for all & GRS for 10–20%'),
10. performing TRS and GRS for all patients ('TRS & GRS for all').

The strategies of interest 1 - 6 were modeled using a dynamic programming, and all allocations of tests and treatments were optimised. These optimised strategies were compared to the reference strategies 7 - 10. The strategies were evaluated based on expected net monetary benefit. Net monetary benefit is defined as

$$\text{NMB} = \text{Health outcomes} \cdot \lambda - \text{Costs} \tag{1}$$

where the health outcomes are measured in QALYs and λ is the societal willingness-to-pay threshold. In this study, the willingness-to-pay threshold is assumed to be 50,000 €/QALY. The net monetary benefit was evaluated over a 10-year time horizon. Hynninen et al. (2019) found the 'TRS & GRS optimized' strategy to be the optimal strategy. In the optimal strategy the 'testing' and 'test-treatment' thresholds for TRS were 10% and 59%, and for GRS they were 17% and 22%. The optimal decision strategy is presented as a flow chart in Figure 2.

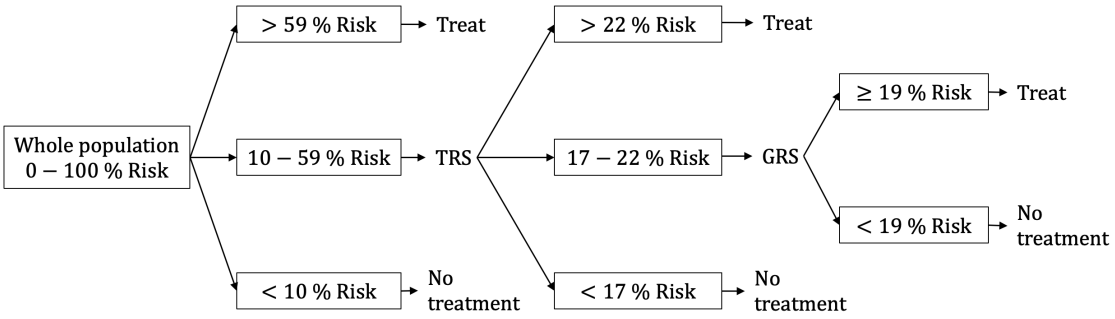

Figure 2: Flow chart representation of the optimal decision strategy from the study by Hynninen et al. (2019). Figure adapted from Hynninen et al. (2019).

3 Methodology

3.1 Decision Programming

Decision Programming is an optimisation framework which has been developed by Salo et al. (2019). The framework combines aspects of stochastic programming and decision analysis. Using this framework, a multi-stage decision problem represented as an influence diagram can be formulated into a mixed integer linear programming (MILP) problem. The strengths of the framework include it being applicable to problems with endogenous uncertainties, meaning that earlier decisions can influence the conditional probabilities of following uncertain events. The framework also allows modeling problems with multiple objectives. The Decision Programming framework has been implemented as a Julia package by Salo et al. (2021). Once a problem is formulated into a MILP using this framework, it can be solved efficiently using off-the-shelf commercial solvers.

The objective of Decision Programming is to solve for an optimal decision strategy Z, which maximises the utility function associated with the problem in question. In the case of this paper, the utility function is the expected net monetary benefit, which is defined in Equation (1). In order to describe how a Decision Programming optimisation problem is structured, several concepts related to influence diagrams, states and paths will be defined in the following.

The problem formulation relies on an influence diagram representation of the problem. The influence diagram is an acyclic graph with nodes $N = C \cup D \cup V$, where C are chance nodes, D are decision nodes and V are value nodes. Each node is represented by a number and the set of numbers is $\{1, ..., |N|\}$, where $|N|$ denotes the number of elements in set N. Dependencies between nodes are denoted by a set of arcs $A = \{(i, j) \mid i, j \in N\}$. The nodes are ordered so that $i < j$ for each arc $(i, j) \in A$ and the value nodes are numbered last, so that $i < j$ for all $i \in C \cup D, j \in V$.

The *information set* $I(j)$ of a node $j \in N$ is the collection of its predecessor nodes. The formal definition of an information set is

$$I(j) = \{i \in C \cup D \mid (i, j) \in A\}. \tag{2}$$

Each chance and decision node $j \in C \cup D$ has a finite set of *states* S_j. A *path* is a sequence of states $s_j \in S_j$ where each chance and decision node $j \in C \cup D$ has a specified state. The *information state* $s_{I(j)}$ of node j is a subpath that includes the states of the information set $I(j)$. The set of all paths is

$$S = \prod_{j \in C \cup D} S_j. \tag{3}$$

For a chance node $j \in C$, the realisation of its state s_j has a stochastic dependence on the realisation of its information state $s_{I(j)}$. To explore this relationship, we define a random variable X_j which describes the realisation of the state of node $j \in C \cup D$. The conditional probability of $X_j = s_j$ is

$$P(X_j = s_j \mid X_{I(j)} = s_{I(j)}) \qquad \forall s_j \in S_j, s_{I(j)} \in S_{I(j)}. \tag{4}$$

For a decision node $j \in D$, the probability of $X_j = s_j$ is dependent on the *local decision strategy* Z_j. A local decision strategy is a function $Z_j : S_{I(j)} \to S_j$, which maps each information state $s_{I(j)}$ to a decision (state) s_j. In Decision Programming, the local decision strategies are modeled as binary variables such that

$$Z_j(s_{I(j)}) = s_j \iff z(s_j \mid s_{I(j)}) = 1 \tag{5}$$

and otherwise $z(s_j \mid s_{I(j)}) = 0$. Note that, since local decision strategies are deterministic functions, $P(X_j = s_j \mid X_{I(j)} = s_{I(j)}) = 1$ whenever $z(s_j \mid s_{I(j)}) = 1$. A *decision strategy* Z contains a local decision strategy for each decision node $j \in D$. A decision strategy Z is said to be *compatible* with a path s if and only if $z(s_j \mid s_{I(j)}) = 1$ for all $j \in D$. The set of all decision strategies is denoted by $\mathbb{Z}$.

The *path probability* of a path $s \in S$ is defined as $\pi(s) = \mathbb{P}(X_{1:|s|} = s_{1:|s|} \mid Z)$, where $|s|$ is the number of elements in path s. The path probability of $s \in S$ is derived by recursively deriving path probabilities of subpaths of s. The recursion is formulated as

$$\pi_k(s) = P(X_k = s_k \mid X_{I(k)} = s_{I(k)})\pi_{k-1}(s), \tag{6}$$

where k is the length of the subpath. The subpaths always start from the first node of path s and the total path probability is $\pi(s) = \pi_{|s|}(s)$. The recursion is initialised by declaring $\pi_0(s) = 1$.

Note that when the kth node in the subpath is a decision node

$$\pi_k(s) = \begin{cases} \pi_{k-1}(s) & \text{if } Z \text{ is compatible with s} \\ 0 & \text{else.} \end{cases} \tag{7}$$

This means that if the decision strategy is not compatible with path s, then $\pi_k(s) = 0$ for some k by the definition of compatibility. Thus, we can define an

upper bound $p(s)$ for the path probability $\pi(s)$ for any $s \in S$ as seen in Equation (8). Notice, that $\pi(s) = p(s)$, when the decision strategy is compatible with path s.

$$p(s) = \prod_{j \in C} P(X_j = s_j \mid X_{I(j)} = s_{I(j)}) \tag{8}$$

Each value node $v \in V$ has a value function $Y_v : S_{I(v)} \to \mathbb{C}$ that maps its information state to a set of consequences $\mathbb{C}$. The utility function $\mathcal{U} : \mathbb{C} \to \mathbb{R}$ maps these consequences to real values. Thus, the utility of a path is defined as the aggregated utility over all consequences of the value nodes

$$\mathcal{U}(s) = \sum_{v \in V} U[Y_v(s_{I(v)})]. \tag{9}$$

In Decision Programming, a problem is formulated as a MILP using the concepts of states, paths and decision strategies. The optimisation problem is formulated as follows

$$\begin{aligned}
\text{max.} \quad & \sum_{s \in S} \pi(s)\mathcal{U}(s) && (10)\\
\text{s.t.} \quad & \sum_{s_j \in S_j} z(s_j \mid s_{I(j)}) = 1, \quad \forall j \in D, s_{I(j)} \in S_{I(j)} && (11)\\
& 0 \leq \pi(s) \leq p(s), \quad \forall s \in S && (12)\\
& \pi(s) \leq z(s_j \mid s_{I(j)}), \quad \forall s \in S && (13)\\
& \pi(s) \geq p(s) + \sum_{j \in D} z(s_j \mid s_{I(j)}) - |D|, \quad \forall s \in S && (14)\\
& z(s_j \mid s_{I(j)}) \in \{0, 1\}, \quad \forall j \in D, s_j \in S_j, s_{I(j)} \in S_{I(j)}, && (15)
\end{aligned}$$

where constraint (11) constrains each local decision strategy to map each information state to exactly one decision. Constraint (12) limits the path probability between zero and the upper bound $p(s)$, which is defined in Equation (8). Constraint (13) ensures that the path probability is zero if the path is incompatible with the decision strategy. Constraint (14) is called the hard lower bound constraint, which ensures that $\pi(s) = p(s)$ if the decision strategy is compatible with the path s, and otherwise the path probability will be zero by constraint (13). The hard lower bound constraint is optional for problems where path utilities $\mathcal{U}(s) \geq 0$ for all $s \in S$. It is needed in maximisation problems where some $\mathcal{U}(s) < 0$ because it is not enforced that the sum of all path probabilities must be equal to one and the paths s for which $\pi(s) > 0$ and $\mathcal{U}(s) < 0$ decrease the objective value. Including the hard lower bound constraint may also enhance solver performance in some problems where $\mathcal{U}(s) \geq 0$ for all $s \in S$. Constraint (15) declares the $z(s_j \mid s_{I(j)})$ variables to be binary variables. Notice that the solver solves for the path probability variables $\pi(s)$ and the decision variables $z(s_j \mid s_{I(j)})$.

Notice that if the upper bound $p(s)$ is zero for a path $s \in S$, then the path probability $\pi(s)$ is also zero according to constraint (12). In this case, the constraints (12)-(14) become trivial and the path probability $\pi(s) = 0$ does not affect the objective

function value. Therefore, in the problem implementation path probability variables for which $p(s) = 0$ are left out of the problem. Naturally, also the constraints associated with these variables are not declared. By leaving out redundant variables and constraints, the size of the model becomes smaller, making it less computationally demanding.

In some decision problems, there are decision combinations which are forbidden due to real world restrictions. This is for example the case in project portfolio selection problems as described by Gustafsson and Salo (2005). The objective in project portfolio selection problems is to manage research and development projects optimally. This entails a multi-stage decision making process where decisions are made about initialising and continuing projects. Naturally, there is a restriction that projects which have not been started cannot be continued. Therefore, the decision combination of continuing a project after not starting it is forbidden. In Decision Programming, forbidden decision combinations are implemented using additional constraints to the problem outlined in Equations (10)-(15). The additional constraint is defined such that

$$\pi(s) \leq 0, \quad \forall s \in F \tag{16}$$

where F is the set of paths that include forbidden decision strategies. This constraint, combined with constraint (12), forces the path probabilities of forbidden paths to zero.

3.2 Constructing the model

The implemented model determines an optimal decision strategy for allocating preventative care for CHD. The model adopts the perspective of the national health care system. The data and structure of the problem are the same as what were used in the study by Hynninen et al. (2019). However, due to the flexibility of Decision Programming, the strategies (1)-(6) mentioned in the literature review do not need to be explicitly defined in the model. This is because all of these strategies are within the feasible solutions of the model and the algorithm solves for the optimal strategy.

The problem setting is such that the patient is assumed to have a prior risk estimate. A risk estimate is a prediction of the patient's chance of having a CHD event in the next ten years. The risk estimates are grouped into risk levels, which range from 0% to 100%. The first testing decision is made based on the prior risk estimate. The first testing decision entails deciding whether TRS or GRS should be performed or if no testing is needed. If a test is conducted, the risk estimate is updated and based on the new information, the second testing decision is made. The second testing decision entails deciding whether further testing should be conducted or not. The second testing decision is constrained so that the same test which was conducted in the first stage cannot be repeated. If a second test is conducted, the risk estimate is updated again. The treatment decision – dictating whether the patient receives statin therapy or not – is made based on the resulting risk estimate of this testing process. Note that if no tests are conducted, the treatment decision is made based on the prior risk estimate.

The influence diagram representation of this problem setting is seen in Figure 3. The orange circular nodes $H, R0, R1, R2$ are chance nodes, which represent the patient's health and the risk levels of their prior and updated risk estimates. Node H represents the uncertainty of whether the patient has a CHD event or remains healthy during the 10 year time frame. Node $R0$ represents the prior risk level before any tests are performed. Node H has the prior risk level $R0$ in its information set, because in the model, we assume that the prior risk accurately describes the probability of having a CHD event. The nodes $R1$ and $R2$ represent the updated risk level after the first and second testing decisions, respectively. If a test is conducted, the risk estimate is updated using Bayesian posterior probability:

$$\textit{Risk estimate} = P(\text{CHD} \mid \text{test result}) = \frac{P(\text{test result} \mid \text{CHD}) \cdot P(\text{CHD})}{P(\text{test result})} \tag{17}$$

where the conditional probabilities $P(\text{test result} \mid \text{CHD})$ are from the study by Abraham et al. (2016) and the probability of having a CHD event, denoted by $P(\text{CHD})$, is the prior risk level $R0$ or the updated risk level $R1$, depending on whether it is the first or second test in question. The denominator $P(\text{test result})$ is calculated as a sum of the numerator and $P(\text{test result} \mid \text{no CHD}) \cdot P(\text{no CHD})$. As the states of nodes R represent the risk levels, the probability of a state in these nodes is the probability of the test updating the risk estimate to that level.

In the influence diagram in Figure 3, decision nodes are illustrated by blue squares. The first and second testing decisions are represented by $T1$ and $T2$, respectively. These decisions determine whether TRS, GRS or no tests should be performed for the patient. Conducting the same test twice is forbidden. Therefore, all paths where the same test is repeated in $T1$ and $T2$ are included in the set of forbidden paths. Furthermore, the forbidden paths include all paths where the first testing decision $T1$ is to not perform testing but then the second testing decision $T2$ is to perform a test. This is because the information yielded from performing only one test is not affected by whether the test is performed in the first or second stage of testing. Therefore, forbidding the paths where no test is performed in $T1$ and TRS or GRS is performed in $T2$ reduces redundancy in the model without information loss. These forbidden path dependencies between nodes $T1$ and $T2$ are not represented by an arrow in the influence diagram because the dependence is of a different nature than the other dependencies represented by arrows.

The final treatment decision is represented by node TD, where the options are to provide treatment or withhold treatment. The treatment decision is made based on the updated risk estimate represented by node $R2$.

There are two value nodes which are used to evaluate the objective function of the model. The value nodes are denoted by green diamonds in the influence diagram. Node TC represents the testing costs. The value function of node TC maps the tests performed – TRS, GRS, TRS & GRS, no tests – to the costs of the tests. The final node of the influence diagram is node HB, which represents the health benefits achieved by a given strategy. The information set of node HB includes the nodes H and TD. The possible information states of node HB are CHD & treatment, CHD

& no treatment, no CHD & treatment, no CHD & no treatment. The node HB maps these strategies to health benefit values. The testing costs and health benefit values were evaluated in the study by Hynninen et al. (2019). The objective function is the expected net monetary benefit and it is evaluated using the testing costs and health benefit values as is shown in Equation (1).

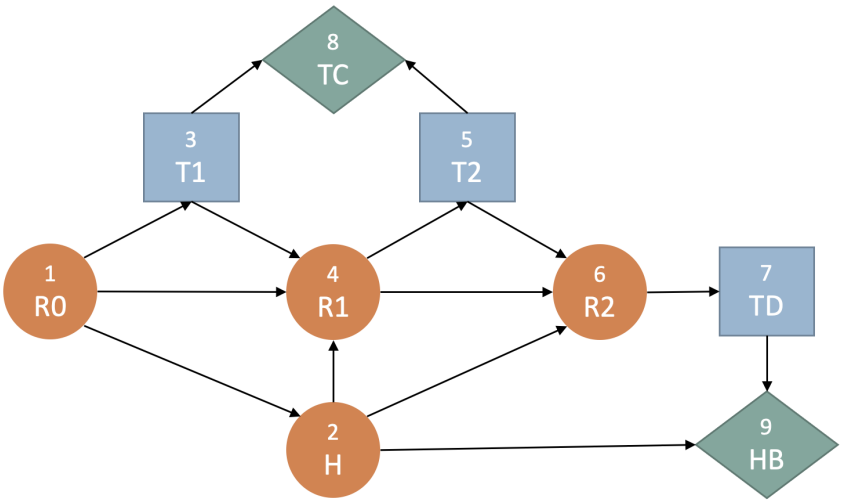

Figure 3: Influence digram representation of the problem.

The states of the nodes in the influence diagram are summarised in Table 1. In the model, the range of risk estimates has been discretised into 101 risk levels $[0\%, 1\%), [1\%, 2\%), ..., [99\%, 100\%), [100\%]$. For brevity, the states of nodes $R0, R1$ and $R2$, representing these risk levels, are referred to by the closed end of the discretised interval. For example, state $0\% = [0\%, 1\%)$. Note that the probability of a state of an R node $P(X_R = s_R)$ describes the chance of the patient being at the risk level represented by that state, not their probability of having a CHD event. For instance, $P(X_{R0} = 1\%)$ describes the probability of an arbitrarily chosen individual having a 1% risk of having a CHD event in the next ten years. In fact, the probabilities of the states of the prior risk node $R0$ describe the population prior risk distribution.

The testing decision nodes $T1$ and $T2$ have states corresponding to the testing options: TRS, GRS, no test. The treatment decision node TD has states representing the choices of providing treatment and withholding treatment. Notice, that in the decision nodes $T1$, $T2$ and TD a separate local decision strategy is derived for each information state 0%, 1%, ..., 100%.

The states of the node H describe whether or not the patient suffers a CHD event in the 10 year time frame. For the value nodes, the states of TC correspond to the possible testing combination and the states of HB correspond to the combinations of health states and treatment options, which determine health benefit outcomes.

The model described above was implemented in the Julia language using the

DecisionProgramming.jl package. The Gurobi Optimization (2021) solver was used for the optimisation.

Table 1: The states S_j for each node $j \in C \cup D \cup V$ in the influence diagram presentation of the problem.

Node	States
R0	{0%, 1%, 2%, ..., 99%, 100%}
H	{CHD, no CHD}
T1	{TRS, GRS, no test}
R1	{0%, 1%, 2%, ..., 99%, 100%}
T2	{TRS, GRS, no test}
R2	{0%, 1%, 2%, ..., 99%, 100%}
TD	{treatment, no treatment}
TC	{TRS, GRS, TRS & GRS, no tests}
HB	{CHD & treatment, CHD & no treatment, no CHD & treatment, no CHD & no treatment}

3.2.1 Model inputs

The model parameters for this problem were derived in the studies by Hynninen et al. (2019) and Abraham et al. (2016). Specifically, the values they derived for the population risk distribution, the conditional probabilities of the TRS and GRS test results, the costs of tests and the health benefit outcomes were used.

The probabilities of the states of node $R0$, which represent the prior risk levels, were set in accordance with the population risk distribution. Setting the probabilities of nodes $R1$ and $R2$ required calculating the probabilities of moving from one risk level to another when a given test is performed. In order to do this, the updated risk estimates were calculated for all possible test results of a given test, for a given prior risk level and state of health – the two health states are that the patient either has a CHD event or not. Equation (17) shows the formula for this update for a single test result when the patient is assumed to have a CHD event. These risk estimates were grouped into risk levels 0%,..., 100% and the probabilities of a risk estimate being updated to each risk level was calculated. The probabilities of these risk levels were used as the probabilities of the states of nodes $R1$ and $R2$ when the information state corresponds to the given prior risk level, testing decision and state of health.

The values of node TC, representing the testing costs, were inserted into the model as QALYs, so that they had the same unit as the health benefits. The testing costs incurred by the possible testing strategies are seen in Table 2. The health benefits corresponding to each information state of node HB are seen in Table 3.

Table 2: The values corresponding to the possible information states of node TC. The testing costs were evaluated by Hynninen et al. (2019).

Testing strategy	Value (QALY)	Value (€)
TRS	$-3.4645 \cdot 10^{-3}$	-173.225
GRS	$-4.0000 \cdot 10^{-3}$	-200
TRS & GRS	$-7.4645 \cdot 10^{-3}$	-373.225
no tests	0	0

Table 3: The values corresponding to the possible information states of node HB. The health benefits were evaluated by Hynninen et al. (2019).

Health and treatment strategy	Value (QALY)
CHD & treatment	6.90
CHD & no treatment	6.65
no CHD & treatment	7.65
no CHD & no treatment	7.70

3.2.2 Model simplification and adjustment

The model described above was found to be computationally intractable due the over 450 000 paths with nonzero probabilities, a large number of which had very small probabilities. To solve the first issue, the model was reduced so that it modeled the problem for only one prior risk level at a time. This modification was executed by changing the chance node $R0$ to a deterministic node. In practice, this meant setting the probability of all states of node $R0$ to zero, with the exception of the chosen prior risk level i:

$$P(X_{R0} = s_{R0}) = \begin{cases} 0 & \forall s_{R0} \neq i \\ 1 & s_{R0} = i. \end{cases} \tag{18}$$

This reduced the number of paths with nonzero path probabilities significantly, because it set all paths starting from the 100 states $s_{R0} \neq i$ to zero. For example, in the model with the prior risk level set to 9%, there were only 10 778 paths with nonzero path probabilities. The model was run for all prior risk levels $i = 0\%, 1\%, ..., 100\%$ in order to find the optimal decision strategy for each subproblem. These strategies were analysed and combined into a general strategy for the full original problem. The objective values of the subproblems were weighted according to the population risk distribution in order to obtain the expected net monetary benefit for the full problem.

The second issue was that approximately half of the nonzero path probabilities were very small (<1e-6). This caused numerical issues in the solver and the problem did not converge to an optimal solution. To mitigate this issue, the path probabilities were scaled by a factor $\alpha > 0$. In practice, this meant scaling the constraints (12)-(14) of the optimisation problem by α as seen in Equations (19)-(21). The value of α was

set to 100, 1000 or 10000, depending on the prior risk level for which the model was being solved.

$$0 \leq \pi(s) \leq \alpha \; p(s), \quad \forall s \in S \tag{19}$$

$$\pi(s) \leq \alpha \; z(s_j \mid s_{I(j)}), \quad \forall s \in S \tag{20}$$

$$\pi(s) \geq \alpha \; (\; p(s) + \sum_{j \in D} z(s_j \mid s_{I(j)}) - |D| \;), \quad \forall s \in S \tag{21}$$

$$\tag{22}$$

4 Results

An optimal decision strategy for allocating preventative care was found using the Decision Programming framework. The results are in line with the results found by Hynninen et al. (2019) using dynamic programming. However, small differences are found in the 'testing', 'test-treatment' and treatment thresholds between the results. Furthermore, the expected monetary benefit found using Decision Programming is larger by a small margin. Notice that the notation introduced in the Section 3.2. will be used in the following. Specifically, the intervals $[0\%, 1\%), [1\%, 2\%), ..., [100\%]$ will be referred to by risk levels 0%, 1%,..., 100%.

The optimal strategy is to perform a two stage screening process using the TRS and GRS testing methods. The optimal strategy in the first testing stage, represented by node $T1$, is to perform TRS testing for patients with a prior risk between 8% and 60%. No testing should be performed for patients with a prior risk less than 8% or above 60%. The optimal strategy in the second testing stage, represented by node $T2$, is to perform further GRS testing for patients with updated risk estimates greater than 15% and less than or equal to 21%. No further testing should be performed for patients with updated risk estimates of less than 15% or greater than 21%. These results are summarised in Table 4.

It is noticed from Table 4 that for the updated risk level $s_{R1} = 15\%$, the optimal strategy in the second testing stage is inconsistent between the described subproblems. This means that the updated risk estimates, which are grouped into risk level 15% in node $R1$ do not lead to the same testing decision in node $T2$. For some of the risk estimates the optimal strategy is to perform GRS and for some to not perform further testing. This inconsistency results from the decision strategies across the 101 subproblems not being interrelated. To obtain a more precise 'testing' threshold for GRS the updated risk estimates, which lead to the inconsistency, were examined. These risk estimates, the prior risk levels s_{R0} before TRS testing and the optimal strategies for $T2$ are shown in Table 5. It is seen that the largest risk estimate, for which no testing is performed is 15.29% and the smallest risk estimate, for which GRS is performed is 15.48%. Therefore, it is concluded that the 'testing' threshold for GRS is between these values. Since, it is not possible to determine the exact threshold, the mean of these two values 15.39% will be reported as the threshold.

Table 4: The optimal local decision strategies for decision nodes $T1$, $T2$ and TD.

Information state	T1	T2	TD
0% - 7%	no test	no test	no treatment
8% - 14%	TRS	no test	no treatment
15%	TRS	no test & GRS	no treatment
16% - 17%	TRS	GRS	no treatment
18%	TRS	GRS	no treatment & treatment
19% - 21%	TRS	GRS	treatment
22% - 60 %	TRS	no test	treatment
61% - 100 %	no test	no test	treatment

Table 5: The risk estimates for which the optimal testing decision in node $T2$ is inconsistent.

Information state s_{R0}	Risk estimate after TRS	Optimal strategy in $T2$
20%	15.02%	no test
11%	15.21%	no test
40%	15.29%	no test
10%	15.48%	GRS
15%	15.74%	GRS
21%	15.82%	GRS
41%	15.84%	GRS
18%	15.90%	GRS

The optimal strategy for the treatment decision, represented by node TD, is to not treat patients with risk estimates below 18% and to treat patients with risk estimates above 18%. The decision strategy is inconsistent for the risk level 18%. This is due to the same reason that was discussed above in the case of testing decision $T2$ for risk level 15%. However, in this case the exact threshold value can be extrapolated based on the health benefit outcomes of treatment strategies and health states. The threshold is found at the point where the expected health benefit is unaffected by whether the patient is treated or not. The expected health benefits with and without treatment are

$$\mathrm{E[treatment]} = p \cdot \mathrm{HB}_{\text{ T \& CHD}} + (1-p) \cdot \mathrm{HB}_{\text{ T \& no CHD}} \tag{23}$$

$$\mathrm{E[no\ treatment]} = p \cdot \mathrm{HB}_{\text{ no T \& CHD}} + (1-p) \cdot \mathrm{HB}_{\text{ no T \& no CHD}} \tag{24}$$

where p is the probability value of the treatment threshold and HB are health benefit values corresponding to different treatment strategies and health states. For conciseness, treatment has been abbreviated with 'T' and no treatment with 'no T' in the indexing of the health benefit values. The health benefit values are found in Table 3. Setting the Equations (23) and (24) equal to each other and solving for the

threshold p gives Equation (25). Using the values found in Table 3, we find that the treatment threshold is $p = 18.63\%$.

$$p = \frac{\mathrm{HB}_{\text{no T \& no CHD}} - \mathrm{HB}_{\text{T \& no CHD}}}{\mathrm{HB}_{\text{T \& CHD}} - \mathrm{HB}_{\text{T \& no CHD}} - \mathrm{HB}_{\text{no T \& CHD}} + \mathrm{HB}_{\text{no T \& no CHD}}}. \tag{25}$$

Based on the results in Table 4 and the analysis of the inconsistent strategies, the optimal strategy can be summarised in a flow chart. The flow chart is shown in Figure 4.

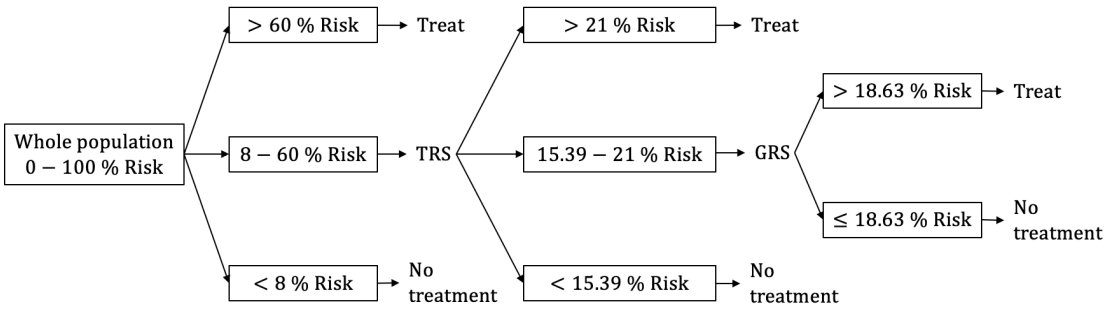

Figure 4: Flow chart representation of the optimal decision strategy for allocating preventative care of CHD.

The optimal strategy seen in Figure 4 is very similar to the strategy found by Hynninen et al. (2019) in Figure 2. The threshold values differ in the strategies by at most 2 risk levels. The strategy found using Decision Programming suggests testing more patients than the strategy found using dynamic programming. Furthermore, the Decision Programming solution suggests treating 0.37% more patients. Despite the Decision Programming strategy administering more tests and more treatment, the expected NMB from the strategy is higher. The threshold values and the expected NMB from both models are show in Table 6. The objective value is shown in QALYs and in euros merely for convenience.

The differences in the results are most probably due to how the range of risk estimates was discretised, the inconsistent strategies among the 101 subproblems and numerical issues. In the dynamic programming model, the range of risk estimates $[0\%, 100\%]$ was discretised by rounding the values to the nearest percentage. However, in the Decision Programming model, the values were grouped into risk levels by rounding down to the nearest percent as described in Section 3.2. The different methods of discretisation may have affected the threshold values for the testing and treatment decisions because the underlying groups of risk estimates used to determine the threshold values were different.

Furthermore, since the Decision Programming model was solved as 101 smaller models, inconsistencies were noticed in the optimal decision strategies of the smaller

models. This was the case when the information state of testing decision node $T2$ was 15% and when the information state of treatment decision node TD was 18%. Resolving these inconsistencies lead to the optimal strategy of the Decision Programming model being more intricate compared the strategy found using dynamic programming. This may partly explain why the Decision Programming model produced a higher expected NMB.

Finally, a challenge in Decision Programming, is that the path probabilities may become very small. Due to this, the integer feasibility tolerance of the solver and the scale factor α had to be adjusted for each model in order for the solution to converge. This leads to issues with numerical accuracy and may cause numerical inconsistencies between the results found in this paper compared to the ones found by Hynninen et al. (2019).

Table 6: The results of the Decision Programming model and of the Dynamic Programming model by Hynninen et al. (2019).

	Decision Programming	Dynamic Programming
TRS 'testing' threshold	8 %	10 %
TRS 'test-treatment' threshold	60%	59 %
GRS 'testing' threshold	15.39%	17 %
GRS 'test-treatment' threshold	22%	22%
Treatment threshold	18.63%	19 %
Objective (QALY)	7.59579	7.59572
NMB (€)	379 790	379 786

5 Discussion and conclusions

An optimal decision strategy for allocating preventative care of CHD was successfully obtained using the Decision Programming framework. The fact that problems with endogenous uncertainty can be modeled using Decision Programming was valuable when modeling this specific problem. This is because the updated risk estimates of a patient are affected by the testing decisions. Furthermore, Decision Programming being applicable to multi-stage decision problems allowed the problem to be modeled in an intuitive way. The testing decisions were represented by two decision nodes, which enabled all possible testing strategies using TRS, GRS or no tests, to be modeled in a way that corresponds to the real life setting of the problem. The third decision node represented the treatment decision, which is made based on the updated risk estimate evaluated in the testing process. Lastly, it was useful that Decision Programming can be applied to problems with several value nodes. In this specific problem, the two value nodes represented the two outcomes that determine NMB – the testing costs and the health benefits. Since these two measures were affected by different decisions – the testing costs were determined by the testing decisions and the health benefits by the treatment decision – it was useful and intuitive that they could be represented by two separate value nodes.

A significant strength of Decision Programming compared to other methods used in cost-benefit analyses is that the explored decision strategies do not need to be predetermined. Furthermore, the decision strategy is optimised using stochastic programming instead of simulation. In earlier studies, similar problems have been modeled using Markov models simulating a set of predefined strategies. In Decision Programming, given the decision options for each decision node, the model generates all possible decision strategies as paths. Then the best feasible strategy is found using linear programming. This allows more strategies to be explored simultaneously. This was especially valuable when modeling this specific problem because the threshold values for testing and treatment decisions are not fixed before the optimisation. Notice that despite this flexibility in the method, the Decision Programming solution was in line with the two-stage screening strategy developed by Hynninen et al. (2019) using dynamic programming.

Some challenges arose when the Decision Programming framework was used to optimise this specific problem. A challenge associated with exploring all possible strategies simultaneously was that the number of paths was very large and thus, the number of path probability variables made the model computationally demanding. The number of paths was large because it is the product of the states over all nodes as seen in Equation (3). Therefore, having nodes with a large number of states – for instance the R nodes with 101 states each – makes the number of paths very large, in this case over 37 million. The number of decision variables was also large because the R nodes, representing the prior and updated risk estimates, formed the information sets of the decision nodes. As a result, the problem was very large and computationally intractable. Therefore, the model was simplified into smaller subproblems by making the prior risk node $R0$ deterministic. This meant that node $R0$ represented only one given risk level and therefore, the model represented the problem for an individual patient instead of the whole population. This modification did not change the number of paths that the subproblems considered when generating the path probability variables. However, since a significantly larger part of the path probabilities were zero, the solver was able to find the optimal solution. Of course, working with 101 subproblems created a challenge in itself, because the parameters of each of the 101 models had to be configured separately and the computing time for many of the subproblems was considerable. In addition, the results of the subproblems had to be processed to find the optimal decision strategy for the original problem.

Another challenge, which arose was that a considerable number of the nonzero path probabilities were very small and caused issues for the solver. The path probabilities are small because they correspond to the product of the probabilities of the states in the path as seen in Equation (8). In this problem, the probabilities of many states were small, because they described the distribution of risk on the range 0%-100%, where some extreme values were very unlikely. In response to this issue, the solver parameters such as integer feasibility tolerance were adjusted. This however did not solve the problem. Thus, the Decision Programming Julia package was modified so that the probabilities could be scaled by a factor α. Scaling the path probabilities by values $\alpha \in \{100, 1000, 10000\}$ increased them so that the solver was able to find a solution.

A possible improvement is reducing the number of paths by removing the deterministic $R0$ node from the models of the subproblems. This would decrease the number of paths in the model by a factor of 101. This way, adding the variables to the model would take considerably less time. This would also reduce the degeneracy in the model, because there would be less paths with zero path probabilities for which the decisions have no effect on the objective function value. Making this modification was decided against because the model would have been a less intuitive representation of the real life problem as the prior risk estimate would not have been represented in it. This modification would not have significantly affected the solution time, because current solvers have presolving capabilities that deal with variables set to zero.

This specific problem could also be explored as a multi-objective optimisation problem using Decision Programming. The two objectives would be reducing testing costs and gaining health benefits. In Decision Programming multiple objectives are represented with multiple value nodes. Since the testing costs and health benefits are already represented as separate value nodes, the problem's influence diagram representation would not need to be modified. With this problem structure, the pareto optimal strategies can be generated using multi-objective optimisation methods. In this specific problem, a strategy Z is pareto optimal if no other feasible strategy has smaller testing costs and greater or equal health benefits, or greater health benefits and smaller or equal testing costs. It would be interesting to explore possible other pareto optimal strategies of this problem, because given the context of the problem, a more preferred strategy may be found. For example, if another pareto optimal strategy exists where greater health benefits are achieved by administering more tests this could be considered a better strategy. This is because people tend to prefer having more tests conducted as opposed to taking unnecessary medication. Furthermore, the adherence to prescribed medication may increase if the patient feels that the medication is more necessary as a result of the additional testing.

In conclusion, the Decision Programming framework was successfully applied in optimising a decision strategy for allocating preventative care for CHD. The properties of the framework allowed the problem to be modeled in a way that was well rooted in its real life setting. Its connection to the influence diagram representation of the problem makes the framework intuitive. The model was found to be computationally demanding, but breaking it down into subproblems and making some modifications to the optimisation problem formulation made it possible to obtain the solution.

References

Gad Abraham, Aki S Havulinna, Oneil G Bhalala, Sean G Byars, Alysha M De Livera, Laxman Yetukuri, Emmi Tikkanen, Markus Perola, Heribert Schunkert, Eric J Sijbrands, et al. Genomic prediction of coronary heart disease. *European heart journal*, 37(43):3267–3278, 2016.

LLC Gurobi Optimization. Gurobi optimizer reference manual, 2021. URL http://www.gurobi.com.

Janne Gustafsson and Ahti Salo. Contingent portfolio programming for the management of risky projects. *Operations research*, 53(6):946–956, 2005.

Yrjänä Hynninen, Miika Linna, and Eeva Vilkkumaa. Value of genetic testing in the prevention of coronary heart disease events. *PLOS ONE*, 14(1):1–16, 01 2019. doi: 10.1371/journal.pone.0210010. URL https://doi.org/10.1371/journal.pone.0210010.

Jamie Jarmul, Mark J. Pletcher, Kristen Hassmiller Lich, Stephanie B. Wheeler, Morris Weinberger, Christy L. Avery, Daniel E. Jonas, Stephanie Earnshaw, and Michael Pignone. Cardiovascular genetic risk testing for targeting statin therapy in the primary prevention of atherosclerotic cardiovascular disease. *Circulation: Cardiovascular Quality and Outcomes*, 11(4):e004171, 2018. doi: 10.1161/CIRCOUTCOMES.117.004171. URL https://www.ahajournals.org/doi/abs/10.1161/CIRCOUTCOMES.117.004171.

M.D. Pauker, Stephen G. and M.D. Kassirer, Jerome P. The threshold approach to clinical decision making. *The New England journal of medicine*, 302(20):1109–1117, May 15 1980. URL https://www.proquest.com/scholarly-journals/threshold-approach-clinical-decision-making/docview/1868774560/se-2?accountid=27468. Copyright - Copyright Massachusetts Medical Society May 15, 1980; Last updated - 2017-12-13.

Eric T. Roberts, Aaron Horne, Seth S. Martin, Michael J. Blaha, Ron Blankstein, Matthew J. Budoff, Christopher Sibley, Joseph F. Polak, Kevin D. Frick, Roger S. Blumenthal, and Khurram Nasir. Cost-effectiveness of coronary artery calcium testing for coronary heart and cardiovascular disease risk prediction to guide statin allocation: The multi-ethnic study of atherosclerosis (mesa). *PLOS ONE*, 10(3):1–20, 03 2015. doi: 10.1371/journal.pone.0116377. URL https://doi.org/10.1371/journal.pone.0116377.

Ahti Salo, Juho Andelmin, and Fabricio Oliveira. Decision programming for multi-stage optimization under uncertainty. *arXiv preprint arXiv:1910.09196*, 2019.

Ahti Salo, Fabricio Oliveira, Juho Andelmin, Olli Herrala, and Jaan Tollander de Balsch. DecisionProgramming.jl: a Julia package for decision programming, 2021. Accessed 20.01.2021. https://gamma-opt.github.io/DecisionProgramming.jl/dev/.

THL. Sydän- ja verisuonitautien yleisyys. *Terveyden ja hyvinvoinnin laitos*, 2020. Updated 03.06.2020. Accessed 21.10.2020. Available: https://thl.fi/fi/web/kansantaudit/sydan-ja-verisuonitaudit/sydan-ja-verisuonitautien-yleisyys.

Emmi Tikkanen, Aki S. Havulinna, Aarno Palotie, Veikko Salomaa, and Samuli Ripatti. Genetic risk prediction and a 2-stage risk screening strategy for coronary heart disease. *Arteriosclerosis, Thrombosis, and Vascular Biology*, 33(9):2261–2266, 2013. doi: 10.1161/ATVBAHA.112.301120. URL https://www.ahajournals.org/doi/abs/10.1161/ATVBAHA.112.301120.

Arnold M. Weissler. Traditional Risk Factors for Coronary Heart Disease. *JAMA*, 291(3):299–299, 01 2004. ISSN 0098-7484. doi: 10.1001/jama.291.3.299-c. URL https://doi.org/10.1001/jama.291.3.299-c.

Sarah Wordsworth, José Leal, Edward Blair, Rosa Legood, Kate Thomson, Anneke Seller, Jenny Taylor, and Hugh Watkins. DNA testing for hypertrophic cardiomyopathy: a cost-effectiveness model. *European Heart Journal*, 31(8): 926–935, 03 2010. ISSN 0195-668X. doi: 10.1093/eurheartj/ehq067. URL https://doi.org/10.1093/eurheartj/ehq067.