DECISION SCIENCES
A JOURNAL OF THE DECISION SCIENCES INSTITUTE

Decision Sciences
Volume 52 Number 5
October 2021

Operationalization of Utilitarian and Egalitarian Objectives for Optimal Allocation of Health Care Resources

Yrjänä Hynninen[†]
Systems Analysis Laboratory, Department of Mathematics and Systems Analysis, Aalto University School of Science, P.O.Box 11100, Aalto, 00076, Finland, e-mail: yrjana.hynninen@gmail.com

Eeva Vilkkumaa
Department of Information and Service Management, Aalto University School of Business (EV), P.O.Box 11100, Aalto, 00076, Finland, e-mail: eeva.vilkkumaa@aalto.fi

Ahti Salo
Systems Analysis Laboratory, Department of Mathematics and Systems Analysis, Aalto University School of Science, P.O.Box 11100, Aalto, 00076, Finland, e-mail: ahti.salo@aalto.fi

ABSTRACT

Resources for health care interventions, such as tests and treatments, are limited. This makes it necessary to prioritize patient segments (defined in terms of their risk) by allocating resources so that the expected contribution to the chosen population-level objective is maximized. In this article, we build a model for the optimal allocation of resources in view of two such objectives: maximizing the aggregate health of the population (utilitarian) and limiting differences in the health outcomes for different patient segments (egalitarian). In particular, we build a two-phase optimization model that (i) first uses dynamic programming to determine what testing and treatment strategies maximize the expected health benefits for each patient segment at different cost levels, and (ii) then solves a binary linear programming problem to determine what resources should be given to each segment to maximize the chosen policy-level objective subject to the overall resource constraint. Our model supports the specification of patient segments, the development of optimal testing and treatment strategies within each segment, and the allocation of available resources to these segments so that the policy-objective will be maximized by implementing these strategies. In addition, the model can be used to guide the interpretation of test results and to assess the impacts of new tests and treatments. It also offers insights into the cost of equity by permitting comparisons between the optimal strategies under utilitarian and egalitarian objectives. We illustrate our approach with real data by optimizing the use of traditional risk scores and genetic tests in preventing coronary heart disease events. [Submitted: April 15, 2019. Revised: January 30, 2020. Accepted: March 23, 2020.]

[†] Corresponding author.

Subject Areas: Coronary heart disease prevention, Health equity, OR in health services, Medical decision making, Multi-objective optimization, and Resource allocation.

INTRODUCTION

Because resources for health care are limited, each unit of resources spent on interventions (such as tests or treatments) means that somewhere or sometime, there are no resources for some other health care purpose (Hunink et al., 2001). This makes it necessary to prioritize patient segments so that resources are spent on those segments in which they can be anticipated to yield the greatest expected benefit. Policies for such resource allocation decisions are usually guided by population-level objectives. These objectives are particularly significant in emergencies and humanitarian health care, where the goal is primarily to increase the population-level health rather than that of any particular individual (Blanchet et al., 2013; Lee, Lavieri, & Volk, 2019). The two most common objectives are *utilitarianism*, which maximizes the aggregate health outcome of a population, and *egalitarianism*, which, for example, seeks to minimize health differences by maximizing the welfare of those who are worst off (Rawls, 1971). In general, these objectives call for different resource allocations. Also, while the utilitarian objective tends to dominate standard economic approaches for the evaluation of public health care interventions (Gold, Siegel, Russell, & Weinstein, 1996; Weatherly et al., 2009), there is evidence that suggests that, in the health domain, the egalitarian approach is more acceptable to most people (Yaari & Bar-Hillel, 1984; Nord, Richardson, Street, Kuhse, & Singer, 1995; Dolan & Cookson, 2000; Cuadras-Morató, Pinto-Prades, & Abellán-Perpiñán, 2001). Indeed, the primary goal of many public health care interventions and programs is to reduce health inequalities, which has inspired much research on the equitable allocation of resources and the assessment of population-level equity considerations (see, e.g., Eddy, 1991; Emanuel, 2000; Beauchamp & Childress, 2001; Sabik & Lie, 2008; Hooker & Williams, 2012).

Cost-effectiveness analysis (CEA) is a well-established approach for evaluating health care programs and interventions (Gold et al., 1996; Drummond, Sculpher, Torrance, O'Brien, & Stoddart, 2005). In CEA, the costs of a program or intervention are compared to the expected aggregate health outcome. Then, decision rules for maximizing health effectiveness (i.e., health gain per one unit of resources consumed) are used to guide resource allocation decisions (e.g., Johannesson & Weinstein, 1993; CADTH, 2017; National Institute for Health and Care Excellence, 2018). Within CEA, methods such as decision trees have been used to model the outcomes of different health care decisions (e.g., Ben-Assuli & Leshno, 2013). In order to identify the optimal decisions, methods such as Mathematical Programming (MP) have been used to allocate resources within specific patient groups and health care programs (e.g., Zaric & Brandeau, 2001; Earnshaw et al., 2002; Brandeau, Zaric, & Richter, 2003; Chalabi, Epstein, McKenna, & Claxton, 2008; Cleary, Mooney, & McIntyre, 2010; Demarteau, Breuer, & Standaert, 2012). Recently, increasing attention has been paid to reconciling CEA with equity

considerations (Ubel, DeKay, Baron, & Asch, 1996; Sassi, Archard, & Le Grand, 2001; Weatherly et al., 2009; Johri & Norheim, 2012). For instance, MP methods have been used to incorporate and assess equity concerns in CEA (e.g., Stinnett & Paltiel, 1996; Epstein, Chalabi, Claxton, & Sculpher, 2007). Generalized cost-effectiveness analysis (GCEA) presents an alternative approach (Murray, Evans, Acharya, & Baltussen, 2000; Edejer et al., 2003; Hutubessy, Chisholm, & Edejer, 2003).

Nevertheless, from the viewpoint of identifying the optimal population-level test and treatment strategies, CEA approaches have shortcomings in that they are typically applied to large, predetermined patient segments with fixed test and treatment strategies (Severens, Sonke, Laheij, Verbeek, & De Vries Robbé, 2001). In particular, because patient segments and testing sequences are not optimized *jointly*, these approaches tend to suggest strategies that are suboptimal in that (i) the resources could be reallocated to achieve a better population-level health outcome or (ii) the same health outcome could be attained with less resources.

Testing strategies have been optimized with partially observable Markov decision processes (POMDP) or compartmental models. Ayer, Alagoz, and Stout (2012), for instance, build a POMDP model to optimize patient-specific mammography screening times. Cevik et al. (2018) present a constrained POMDP model to study the optimal allocation of limited mammography resources to screen a population. Lee et al. (2019) optimize the use of limited resources for the screening of a population for hepatocellular carcinoma by modeling the problem as a family of restless bandits in which each patient's disease progression is assumed to evolve as a POMDP. Güneş, Örmeci, and Kunduzcu (2015) present a compartmental model for allocating limited colonoscopy resources between screening and diagnostic services, whereas Deo, Rajaram, Rath, Karmarkar, and Goetz (2015) deploy a compartmental model to plan for HIV screening, testing, and care.

The above approaches are useful when the possible deterioration in a patient's state of health over time needs to be accounted for. Nevertheless, all the above studies assume fixed interpretations of test results (i.e., predetermined positivity thresholds) and fixed treatment decisions based on binary test results (i.e., predetermined treatment thresholds). They also assume that the final screening test is perfect, and that a positive test result is straightforwardly followed by a treatment. In particular, none of these studies recognizes that the optimal interpretation of test results, as well as the optimal treatment decision based on these results, may both depend on how much resources there are.

Our main contribution lies in developing models that guide the optimal allocation of resources between patient segments for both utilitarian and egalitarian population-level objectives, in recognition of the full range of alternative tests and treatments; the available resources; and the distribution of patients in segments representing different risk levels in the population. Toward this end, we build a two-phase optimization model. In the first phase, the optimal test and treatment strategies are determined for all patient segments, defined as nonoverlapping intervals of prior probabilities with which the patient has a given disease. For each segment, all Pareto optimal testing and treatment strategies are determined with dynamic programming by maximizing the expected health outcome subject to bounds on expected costs. In the second phase, the sizes of these patient segments (i.e., their

share of the total population) are accounted for by solving a resource-constrained binary linear programming problem to determine which combinations of these Pareto optimal testing and treatment strategies maximize the chosen population-level objective (utilitarian or egalitarian). We illustrate the decision model by optimizing the use of traditional risk scores and genetic tests in preventing coronary heart disease (CHD) events, based on real data from Finnish national health care registers, the Finnish Institute for Molecular Medicine, and published literature.

The contribution of this article to existing literature is fourfold. First, while earlier CEA evaluations of health care interventions are based on comparing the application of few, predetermined intervention strategies within large predefined patient segments, our model helps optimize these strategies separately for each risk level, thereby supporting the specification of meaningful patient segments. Thus, the model provides defensible policy recommendations on how to allocate resources (i) between different patient segments and (ii) between tests and treatments. Moreover, because our model can be employed to optimize the use of new tests and treatments along with existing ones, it helps assess whether the new ones are cost-effective. Second, unlike earlier approaches, the model supports the optimal interpretation of test results and the optimal selection of treatments for each patient segment and at different resource levels. By doing so, the model helps convert policy decisions about the population-level objective and budget into optimal intervention strategies at the operational level. Third, although earlier approaches often monetize health outcomes through parameters such as the willingness-to-pay (WTP) threshold to support resource allocation decisions, we use multi-objective optimization techniques to generate the entire frontier of Pareto optimal population-level strategies. This approach enables policy-level analyses on, for example, the marginal increase in expected health outcome that could be obtained by increasing the level of resources by some amount. Moreover, it also helps prepare for health care emergencies by suggesting what resources would be needed to reach the population-level objectives satisfactorily in scenarios representing different population distributions over patient segments. Finally, we solve the optimal test and treatment strategies at different cost levels for both the utilitarian and egalitarian objective. This makes it possible to assess the cost of equity, defined as the difference between (i) health outcomes at a given resource level or (ii) the level of resources required to obtain a given health outcome (e.g., Stinnett & Paltiel, 1996; Weatherly et al., 2009).

The rest of this article is structured as follows. The following section develops the model for determining optimal population-level test and treatment strategies subject to resource constraints under utilitarian and egalitarian objectives. Then, an illustrative example on the screening strategy of CHD events is presented. The article concludes with a discussion of the limitations, assumptions, and policy implications of our results.

A DECISION MODEL FOR POPULATION-LEVEL STRATEGIES

We consider a population of patients who belong to segments based on their initial probability of having a given disease. In the two-phase process for optimizing the population-level test and treatment strategy, the first phase gives the Pareto

Figure 1: Two-phase optimization process. First, identify Pareto optimal strategies for each patient segment (left). Pareto optimal strategies are depicted with the symbol • and Pareto dominated strategies with the symbol ○. Second, identify the optimal combination of Pareto optimal strategies for patient segments subject to a cost constraint (right).

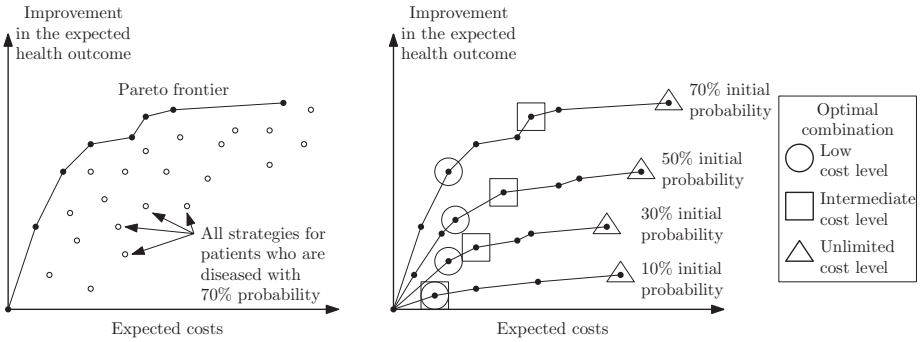

optimal strategies (i.e., strategies for which the expected health outcome cannot be improved without increasing expected costs) for every segment. In the second phase, these Pareto optimal strategies are combined with information about the number of patients in each segment in order to maximize the utilitarian and egalitarian objectives subject to a resource constraint.

Figure 1 shows graphs for these two phases. The left-hand graph illustrates phase 1 for a patient whose initial probability of disease is 70%. Here, each dot represents the expected health outcome and cost of one possible test-treatment strategy. Pareto dominated strategies are marked with white dots and Pareto optimal strategies with black dots. The right-hand graph illustrates phase 2 for three different cost levels ("Low," "Intermediate," and "Unlimited"). The population-level optimal strategy is a combination of Pareto optimal strategies for different patient segments (here, corresponding to 10%, 30%, 50%, and 70% initial probabilities) such that the combination maximizes the population-level objective subject to the population-level resource constraint, taking into account the relative shares of patients in different segments.

Model for Test and Treatment Strategies

States of health and treatment actions

The patient's state of health is modeled as a discrete random variable S with realizations $s \in \{0, 1\}$ so that the patient either does not $(S = 0)$ or does have $(S = 1)$ a given disease. This state is static and does not change during the testing period. Uncertainty about the state is represented by probability $p := p(S = 1)$ so that $p(S = 0) = 1 - p$. Based on the probability of disease p (cf. *risk of disease*; Sox, Higgins, & Owens, 2013), one of the actions $a \in \mathcal{A} = \{0, \ldots, n_A\}$ is selected, whereby $a = 0$ stands for no treatment and $a = 1, \ldots, n_A$ are alternative treatments. The health outcome $h(a|S = s)$ resulting from executing

action $a \in \mathcal{A}$ to a patient in state $s \in \{0, 1\}$ is assumed to be known for all actions and both states. This health outcome can be measured, for instance, in terms of quality-adjusted life-years (QALYs) (e.g., Sox et al., 2013). For brevity, we denote $h(a|1) := h(a|S = 1)$ and $h(a|0) := h(a|S = 0)$.

Treatment costs are considered from the perspective of the health care system and they are assumed to be known. Specifically, $c^{\text{treat}}(a|s)$ denotes the cost of treating a patient in state $s \in \{0, 1\}$ with action a (e.g., medication or a surgery operation), while $c^{\text{post}}(a|s)$ denotes the posttreatment cost for $s \in \{0, 1\}$ and $a \in \mathcal{A}$. For example, the posttreatment cost of not treating a patient with a disease (i.e., $c^{\text{post}}(0|S = 1)$) may include the costs of an acute disease event and 1-year follow-up. For brevity, we denote $c^{\text{post}}(a|1) := c^{\text{post}}(a|S = 1)$ and $c^{\text{post}}(a|0) := c^{\text{post}}(a|S = 0)$.

Tests and test results

Information about the patient's state of health can be obtained through tests $t \in \mathcal{T} = \{0, \ldots, T\}$, where $t = 0$ represents the option of not testing. The cost of test t is $c^{\text{test}}(t)$, where $c^{\text{test}}(0) = 0$. The result of test t is a discrete random variable R_t whose realization r_t belongs to the set of n_t possible test results $\{r_{t,1}, \ldots, r_{t,n_t}\}$. This result depends only on the patient's true state of health s and not on the results of other tests, that is, the tests are assumed to be conditionally independent. We write $p(r_t|s) := p(R_t = r_t|S = s)$ and $p(s|r_t) := p(S = s|R_t = r_t)$.

There are at most K testing stages and the maximum number of tests at stage $k \in \{1, \ldots, K\}$ is $n(k)$. The combination of tests at stage k is denoted by $\boldsymbol{t}^k = (t_1^k, \ldots, t_{n(k)}^k)$. There can be tests that cannot (or should not) be carried out repeatedly, for instance, if these tests do not provide additional information upon repetition (for instance, the results of genetic testing would not change). The tests that can be used multiple times are denoted by $\mathcal{M} \subseteq \mathcal{T}$. Each selected test t_i^k must belong to the set $\mathcal{T}_A^k \subseteq \mathcal{T}$ of available tests at stage k. At the first stage $k = 1$, the set of available tests is $\mathcal{T}_A^1 = \mathcal{T}$, after which this set is updated depending on what previous tests have been carried out. In particular, let us denote by $\mathcal{T}_U^k = \{t \mid \exists\, i \in \{1, \ldots, n(k)\} \text{ s.t. } t_i^k = t\}$ the set of tests that are carried out at stage k. Then, the set of available tests at stage $k + 1$ is updated to $\mathcal{T}_A^{k+1} = \mathcal{T}_A^k \backslash (\mathcal{T}_U^k \backslash \mathcal{M})$, where $(\mathcal{T}_U^k \backslash \mathcal{M}) = \{t | t \in \mathcal{T}_U^k, t \notin \mathcal{M}\}$ contains tests that are carried out at stage k but cannot be repeated later. For all $t_i^k \notin \mathcal{M}$ and all k, constraints $t_i^k \neq t_j^k \; \forall\, i, j = 1, \ldots, n(k), i \neq j$ ensure that tests that cannot be carried out multiple times are not repeated at any stage k. If all tests can be repeated, then $\mathcal{M} = \mathcal{T}$ and the set of available tests stays the same across all stages so that $\mathcal{T}_A^k = \mathcal{T}_A^1 = \mathcal{T} \; \forall\, k$. Figure 2 illustrates the process of updating the set of available tests.

The results of test combination $\boldsymbol{t}^k$ are denoted by $\boldsymbol{R}^k = (R_{t_1^k}, \ldots, R_{t_{n(k)}^k}) \rightarrow \boldsymbol{r}^k = (r_{t_1^k}, \ldots, r_{t_{n(k)}^k})$. The cost of a test combination is assumed to be the sum of the costs of the tests that it contains. An additional testing stage is assumed to have a fixed cost c^{stage}, which could arise from setup costs or the impact of delays, reflecting the possible deterioration of a patient's state of health during the testing period. This fixed cost is incurred only if there is at least one test at stage k, in which case $t_{n(k)}^k > 0$. Hence, the cost of test combination $\boldsymbol{t}^k$ is $c^{\text{test}}(\boldsymbol{t}^k) = \sum_{i=1}^{n(k)} c^{\text{test}}(t_i^k) + c^{\text{stage}}$ if $t_{n(k)}^k \neq 0$, and $c^{\text{test}}(\boldsymbol{t}^k) = 0$ otherwise.

Figure 2: Updating the set of available tests. Black dots represent which of the available tests can be carried out multiple times ($\mathcal{T}_A^k \cap \mathcal{M}$) and circles represent those that can be carried out only once ($\mathcal{T}_A^k \backslash \mathcal{M}$).

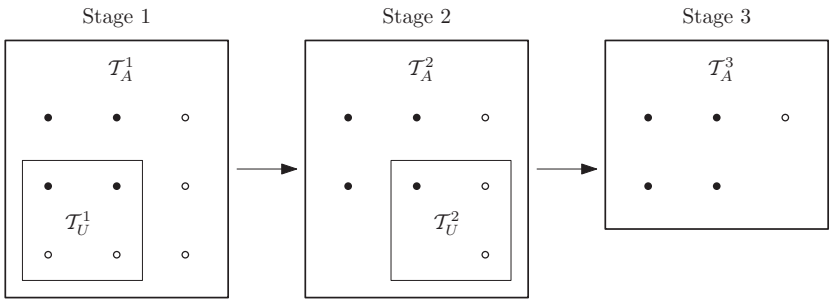

Updating health state probabilities

At each stage k, the probabilities of health states $s \in \{0, 1\}$ are updated based on test results $\boldsymbol{r}^k = (r_{t_1^k}, \ldots, r_{t_{n(k)}^k})$. Specifically, let p_{k-1} denote the *prior probability* that the patient has the disease at the beginning of stage k and assume that test results $\boldsymbol{r}^k$ are obtained with conditional probability $p(\boldsymbol{r}^k|S = s)$ when the patient is in health state $s \in \{0, 1\}$. Then, the *posterior probability* $p_k := p(1|\boldsymbol{r}^k)$ for the event that the patient has the disease is given by Bayes' rule as

$$p_k = \frac{p(\boldsymbol{r}^k|S = 1) \cdot p_{k-1}}{p(\boldsymbol{r}^k|S = 1) \cdot p_{k-1} + p(\boldsymbol{r}^k|S = 0) \cdot (1 - p_{k-1})}. \tag{1}$$

Here, the denominator represents the probability of obtaining the results $\boldsymbol{r}^k$ given prior probability p_{k-1}. For brevity, we denote this probability by $p(\boldsymbol{r}^k|p_{k-1}) := p(\boldsymbol{r}^k|S = 1) \cdot p_{k-1} + p(\boldsymbol{r}^k|S = 0) \cdot (1 - p_{k-1})$.

The posterior probability p_k in (1) from stage k is the prior probability for the next stage $k + 1$. The *initial prior probability* p_0 can be estimated, for instance, based on how prevalent the disease is among people of the same age and gender. It can also be adjusted by using supplementary information such as family history or clinical history.

Phase 1: Solving the Pareto Optimal Strategies for a Single Patient

Pareto optimal testing and treatment strategies

At each stage k of the K-stage decision process, a decision maker (DM) can decide (i) to carry out test combination $\boldsymbol{t}^k$ to obtain information about the patient's state of health or (ii) to stop testing and carry out some action a, including the option $a = 0$ of not treating. A *test and treatment strategy* is a sequence of decisions in which the kth stage decision is informed by all earlier test results. This strategy can be illustrated through paths in a decision tree. Figure 3, for example, shows a decision tree for two testing stages, one treatment stage, and two dichotomized tests that cannot be repeated. The bolded lines show the optimal paths defined by the decisions that maximize the expected health outcome. The assumption that a

Figure 3: A decision tree representing alternative test and treatment strategies. Bolded lines show optimal choices at decision nodes (squares).

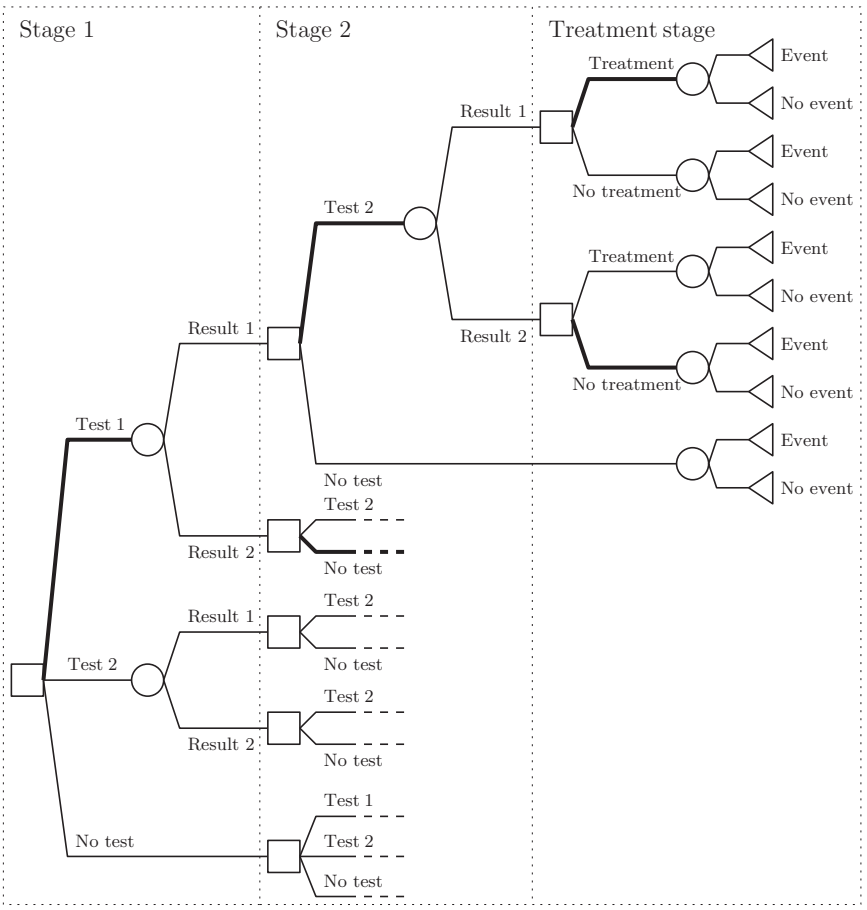

patient goes through all K testing stages is purely technical rather than restrictive, because the treatment decision can be made immediately after any testing stage or even based on the initial disease probability.

For each strategy, the expected cost C and expected health outcome H can be computed for any initial disease probability p_0. A strategy dominates another in the Pareto sense if it offers a higher expected health outcome H at a lower expected cost C, allowing for the possibility that there can be an equality in one (but not both) of these two comparisons.

Definition 1: *Let* (C, H) *be a test and treatment strategy with expected cost* C *and expected health outcome* H. *Strategy* (C, H) Pareto dominates *strategy* (C', H'), *denoted* $(C, H) \succ (C', H')$, *if*

$$C \leq C' \text{ and } H \geq H',$$

where at least one of the inequalities is strict.

Because a rational DM would not choose a Pareto dominated strategy, it is reasonable to focus on Pareto optimal strategies.

Definition 2: *Strategy* (C, H) *is* Pareto optimal*, if* $\nexists$ (C', H') *such that* $(C', H') \succ (C, H)$.

Thus, the expected health outcome of a Pareto optimal strategy cannot be improved without increasing the expected cost, and the expected cost cannot be decreased without decreasing the expected health outcome.

The ε-constraint method for approximating the set of Pareto optimal strategies

In general, the set of optimal strategies is nonconvex in that a testing strategy that is optimal for patient segments p_{i-1} and p_{i+1} at some cost level may not be optimal for an intermediate segment p_i. Hence, analytical results on optimal strategies can be obtained only in simple special cases (see Appendix A for such results on optimal strategies at lowest or highest cost levels for a single test with binary results). In realistic problems involving multiple tests, test results, and testing stages, numerical solution methods must therefore be employed. In this article, we use ε-constraint method (e.g., Laumanns, Thiele, & Zitzler, 2006; Deb, 2014) to approximate the Pareto optimal strategies for a patient with initial disease probability p_0. That is, we first generate a sequence $b_1, \ldots, b_J$ of upper bounds on the expected cost C of the strategy, where $b_j = b_{j-1} + \varepsilon,\ \forall j \in \{2, \ldots, J\}$ for some $\varepsilon > 0$. Then, we solve J single-objective optimization problems to determine which strategies maximize the expected health outcome H subject to every upper bound b_j in this sequence.

The ε-constraint method helps identify also those Pareto optimal strategies that are below the convex hull of the Pareto frontier. Such strategies are fairly common in the context of optimizing test and treatment strategies (see Appendix B for an example) but cannot be found by methods in which the health outcomes are first monetized through some parameter (e.g., by using a WTP threshold), after which the value of this parameter is varied (cf. the weighted sum approach; Das & Dennis, 1997). Strategies below the convex hull of the Pareto frontier would be dominated by those obtained through randomly assigning a share $w \in (0, 1)$ of patients to a strategy on the convex hull with a lower cost and lower health outcome, and the remaining share $(1 - w) \in (0, 1)$ to another strategy on the convex hull with a higher cost and higher health outcome such that the expected cost of this randomized strategy would coincide with that of the original one (cf. extended dominance; Cantor, 1994). However, as such randomized strategies are ethically untenable in health care, the strategies under the convex hull represent viable compromise solutions between costs and health outcomes.

The ε-constraint method characterizes Pareto optimal strategies through an approximation whose goodness depends on ε. The smaller the value of ε, the more Pareto optimal strategies are generated but, on the other hand, the larger the number $J = (b_J - b_1)/\varepsilon + 1$ of single-objective optimization problems that need to be solved. The choice of ε thus involves a trade-off between computation time and the quality of the approximation.

Dynamic programming solution for maximizing the expected health outcome

The strategy that maximizes the expected health outcome for a given initial disease probability p_0 subject to an upper bound b_j on the expected cost can be solved with dynamic programming based on Bellman's principle of optimality (Bellman, 1957). At stage k, the state is defined by (i) the current probability p_{k-1} that the patient has the disease, (ii) the set of available tests $\mathcal{T}_A^k$, and (iii) the remaining resources from stage k onward given probability p_{k-1}, denoted by $b_{k,p_{k-1}}$. At stage 1, the state consists of the initial disease probability p_0, tests $\mathcal{T}_A^1 = \mathcal{T}$, and the selected cost bound $b_{1,p_0} = b_j$.

By Bellman's principle of optimality, we first solve the optimal treatment decisions at stage $k = K + 1$ for all states and then proceed recursively to the earlier stages $k = K, \ldots, 1$ to optimize decisions for all states while accounting for optimal decisions at the later stages. That is, the dynamic programming algorithm first solves for every state $(p_K, \mathcal{T}_A^{K+1}, b_{K+1,p_K})$ the treatment action a that maximizes the expected health outcome at the last stage $k = K + 1$ subject to the amount of available resources for treatment b_{K+1,p_K} given prior probability p_K:

$$H_{K+1}\big(p_K, \mathcal{T}_A^{K+1}, b_{K+1,p_K}\big) = \max_a \overbrace{\big(h(a|1) \cdot p_K + h(a|0) \cdot (1 - p_K)\big)}^{\text{Expected health outcome}}$$

$$\text{subject to} \quad \underbrace{c^{\text{treat}}(a|1) \cdot p_K + c^{\text{treat}}(a|0) \cdot (1 - p_K)}_{\text{Expected treatment cost}}$$

$$+ \underbrace{c^{\text{post}}(a|1) \cdot p_K + c^{\text{post}}(a|0) \cdot (1 - p_K)}_{\text{Expected post-treatment cost}} \leq b_{K+1,p_K}.$$

Denoting by a^* the optimal solution to the above problem, the optimal expected cost for the state $(p_K, \mathcal{T}_A^{K+1}, b_{K+1,p_K})$ is

$$C_{K+1}\big(p_K, \mathcal{T}_A^{K+1}, b_{K+1,p_K}\big) = \big(c^{\text{treat}}(a^*|1) + c^{\text{post}}(a^*|1)\big) \cdot p_K + \big(c^{\text{treat}}(a^*|0) + c^{\text{post}}(a^*|0)\big) \cdot (1 - p_K).$$

In the last testing stage $k = K$, it is necessary to decide for each state $(p_{K-1}, \mathcal{T}_A^K, b_{K,p_{K-1}})$ (i) the optimal level of resources b_{K+1,p_K} left for treatment and posttreatment, corresponding to each potential posterior probability p_K that can be achieved by a given test combination $\boldsymbol{t}^K$ and (ii) the optimal test combination $\boldsymbol{t}^K$ subject to the upper bound $b_{K,p_{K-1}}$ on testing costs at stage $k = K$ and expected treatment and posttreatment costs at stage $k = K + 1$. We use $\boldsymbol{b}_{K+1}$ to denote the vector of resource levels left for treatment and posttreatment with elements b_{K+1,p_K} corresponding to the potential values of posterior probability p_K. Then, the optimal decisions can be found by a solving a two-stage stochastic optimization problem:

$$H_K\big(p_{K-1}, \mathcal{T}_A^K, b_{K,p_{K-1}}\big)$$

$$= \max_{\boldsymbol{t}^K} \max_{\boldsymbol{b}_{K+1}} \overbrace{\sum_{\boldsymbol{r}^K} H_{K+1}\big(p_K, \mathcal{T}_A^{K+1}, b_{K+1,p_K}\big) \times p\big(\boldsymbol{r}^K | p_{K-1}\big)}^{\text{Expected health outcome}}, \tag{2}$$

$$\text{subject to } \underbrace{\sum_{\boldsymbol{r}^K} \left(C_{K+1}\left(p_K, \mathcal{T}_A^{K+1}, b_{K+1,p_K}\right) \times p\left(\boldsymbol{r}^K | p_{K-1}\right)\right)}_{\text{Expected treatment and post-treatment cost}}$$
$$+ \underbrace{c^{\text{test}}(\boldsymbol{t}^K)}_{\text{Testing cost at stage } K} \leq b_{K,p_{K-1}}.$$

In the inner maximization problem, the test combination $\boldsymbol{t}^K$ is fixed but its result $\boldsymbol{R}^K$ is random so that the posterior probability $p_K := p(1|\boldsymbol{R}^K)$ of the patient having the disease is random, too. The problem is to find the optimal resource levels $\boldsymbol{b}_{K+1}$ for treatment stage $k = K+1$ for all possible posterior probabilities p_K such that (i) the expected health outcome is maximized and (ii) the sum of testing costs at stage $k = K$ and expected treatment and posttreatment costs at stage $k = K+1$ is at most $b_{K,p_{K-1}}$. Solving the outer maximization problem gives the optimal test combination $\boldsymbol{t}^K$. Denoting by $(\boldsymbol{t}^{K*}, \boldsymbol{b}_{K+1}^*)$ the optimal solution to problem (2) and by $\boldsymbol{r}^{K*}$ the possible results of test combination $\boldsymbol{t}^{K*}$, the optimal expected cost corresponding to state $(p_{K-1}, \mathcal{T}_A^K, b_{K,p_{K-1}})$ is

$$C_K\left(p_{K-1}, \mathcal{T}_A^K, b_{K,p_{K-1}}\right) = \sum_{\boldsymbol{r}^{K*}} \left(C_{K+1}\left(p_K, \mathcal{T}_A^{K+1}, b_{K+1,p_K}\right) \times p\left(\boldsymbol{r}^{K*} | p_{K-1}\right)\right) + c^{\text{test}}(\boldsymbol{t}^{K*}).$$

The optimal decisions $(\boldsymbol{t}^{k*}, \boldsymbol{b}_{k+1}^*)$ for the preceding testing stages $k = K-1, \ldots, 1$ can be solved similarly. The recursive algorithm for finding optimal testing and treatment strategies is summarized below.

Recursive algorithm for maximizing the expected health outcome:

For treatment stage $k = K+1$:

$$H_{K+1}\left(p_K, \mathcal{T}_A^{K+1}, b_{K+1,p_K}\right) = \max_a \left(h(a|1) \cdot p_K + h(a|0) \cdot (1 - p_K)\right) \tag{3}$$

$$\text{subject to } \left(c^{\text{treat}}(a|1) + c^{\text{post}}(a|1)\right) \cdot p_K + \left(c^{\text{treat}}(a|0) + c^{\text{post}}(a|0)\right) \cdot (1 - p_K) \leq b_{K+1,p_K}. \tag{3}$$

For testing stages $k = K, \ldots, 1$:

$$H_k\left(p_{k-1}, \mathcal{T}_A^k, b_{k,p_{k-1}}\right) = \max_{\boldsymbol{t}^k} \max_{\boldsymbol{b}_{k+1}} \sum_{\boldsymbol{r}^k} H_{k+1}\left(p_k, \mathcal{T}_A^{k+1}, b_{k+1,p_k}\right) \times p\left(\boldsymbol{r}^k | p_{k-1}\right), \tag{4}$$

$$\text{subject to } \sum_{\boldsymbol{r}^k} \left(C_{k+1}\left(p_k, \mathcal{T}_A^{k+1}, b_{k+1,p_k}\right) \times p\left(\boldsymbol{r}^k | p_{k-1}\right)\right) + c^{\text{test}}(\boldsymbol{t}^k) \leq b_{k,p_{k-1}},$$

where

$$C_{K+1}\left(p_K, \mathcal{T}_A^{K+1}, b_{K+1,p_K}\right) = c^{\text{post}}(a^*|1) \cdot p_K + c^{\text{post}}(a^*|0) \cdot (1 - p_K) + c^{\text{treat}}(a^*) \tag{5}$$

$$C_k\left(p_{k-1}, \mathcal{T}_A^k, b_{k,p_{k-1}}\right) = \sum_{\boldsymbol{r}^{k*}} C_{k+1}\left(p_{k+1}, \mathcal{T}_A^{k+1}, b^*_{k+1,p_{k+1}}\right) \times p\left(\boldsymbol{r}^{k*}|p_k\right). \quad (6)$$

At stage $k = 1$, $H_1(p_0, \mathcal{T}, b_j)$ and $C_1(p_0, \mathcal{T}, b_j)$ are the expected health outcome and cost of the optimal test and treatment strategy for a patient with initial disease probability p_0, subject to the cost bound b_j.

Computation of Pareto optimal strategies

The optimization problems (3) and (4) in the recursive algorithm are highly nonlinear in the initial disease probability p_0 due to the repeated use of the Bayes' formula (1). Hence, the establishment of optimality conditions for the set of Pareto optimal strategies—through, for instance, identifying intervals of prior probabilities within which the set of Pareto optimal strategies is the same—is very challenging if not impossible. Consequently, we solve the Pareto optimal strategies for all initial disease probabilities $p_0 \in [0, 1]$ numerically. To do this, we discretize the probability interval [0,1] into I points $p_{k-1} \in \{p^1, \ldots, p^I\}$ such that $p^i = (i-1)/(I-1)$ for all $i \in \{1, \ldots, I\}$. If $I = 101$, for instance, then $p^1 = 0\%$, $p^2 = 1\%, \ldots, p^{101} = 100\%$. The levels of available resources $b_{k,p_{k-1}}$ at stage k for probability p_{k-1} are discretized similarly to the sequence of upper bounds on the expected cost C of the strategy: $b_{k,p_{k-1}} \in \{b_1, b_1 + \varepsilon, \ldots, b_J - \varepsilon, b_J\}$ for all $k = 1, \ldots, K+1$. The recursive algorithm (3)–(6) is then carried out for all combinations of discretized values of p_{k-1} and $b_{k,p_{k-1}}$, $k \in \{1, \ldots, K+1\}$.

The expected health outcome and expected cost for initial disease probability p^i and upper bound b_j on the expected cost of the strategy are $H_1(p^i, \mathcal{T}, b_j)$ and $C_1(p^i, \mathcal{T}, b_j)$, respectively. If an increase from b_{j-1} to b_j in the upper bound on the expected cost does not increase the expected cost of the optimal strategy (i.e., if $C_1(p^i, \mathcal{T}, b_{j-1}) = C_1(p^i, \mathcal{T}, b_j)$), the optimal strategies for the upper bounds b_{j-1} and b_j are the same. The approximated set of Pareto optimal strategies for initial probability p^i consists of the $J_i \leq J$ unique strategies. The expected costs and health outcomes of these strategies are denoted by c_{i,j_i} and h_{i,j_i}, $j_i \in \{1, \ldots, J_i\}$, respectively.

Phase 2: Optimizing the Population-Level Strategy

In the second phase, we account for the distribution of different initial disease probabilities in the population and identify which combination of Pareto optimal testing and treatment strategies maximizes the population-level objective (utilitarian or egalitarian) subject to a resource constraint. More specifically, the population is divided into patient segments $i \in \{1, \ldots, I\}$ that correspond to different discretized values p^i of the patients' initial disease probabilities p_0. These probabilities can be estimated based on, for example, the prevalence of the disease in various subpopulations defined by the patients' age and gender as well as other relevant information.

The number of patients in segment i is d_i. Based on the population distribution d_i over initial disease probabilities p^i, the problem is to determine the optimal combination of segment-specific Pareto optimal strategies $j_i \in \{1, \ldots, J_i\}$,

$i \in \{1, \ldots, I\}$ with expected costs c_{i,j_i} and health outcomes h_{i,j_i} such that the total expected cost of this combination does not exceed some resource level B. Decisions concerning the selection of Pareto optimal strategies are represented by binary decision variables $x_{i,j_i} \in \{0, 1\}$ such that $x_{i,j_i} = 1$ if and only if strategy j_i is carried out to patient segment i. The vector of decision variables corresponding to patient segment i is denoted by $\boldsymbol{x}_i \in \{0, 1\}^{J_i}$, and the vector of all decision variables is denoted by $\boldsymbol{x} \in \{0, 1\}^M$, where $M = \sum_i J_i$.

Utilitarian approach

In the utilitarian approach, the objective is to maximize the population-level health outcome subject to the population-level resource constraint. Computationally, it is implemented by the following binary linear programming problem, referred to as PROBLEM U-ALLOCATION:

$$U^* = \max_{\boldsymbol{x}} \sum_{i=1}^{I} \sum_{j_i=1}^{J_i} x_{i,j_i} d_i h_{i,j_i}$$

$$\text{subject to } \sum_{i=1}^{I} \sum_{j_i=1}^{J_i} x_{i,j_i} d_i c_{i,j_i} \leq B \tag{7}$$

$$\sum_{j=1}^{J_i} x_{i,j_i} = 1 \text{ for all } i \in \{1, \ldots, I\} \tag{8}$$

$$x_{i,j_i} \in \{0, 1\} \text{ for all } i \in \{1, \ldots, I\},\ j_i \in \{1, \ldots, J_i\}, \tag{9}$$

where d_i is the number of patients in segment i corresponding to initial disease probability p^i; h_{i,j_i} and c_{i,j_i} are the expected patient-specific health outcome and cost resulting from the application of strategy j_i to segment i; and B is the population-level resource constraint.

Egalitarian approach

In the egalitarian approach, the objective is to reduce health differences between individuals in the population. In this article, we model egalitarianism by maximizing the health outcome for the worst off. This leads to a maximin binary linear programming problem with objective function $\max_{\boldsymbol{x}} \min_i \ x_{i,j_i} h_{i,j_i}$, and constraints (7)–(9). A model with a maximin objective function would attempt to allocate resources equitably among segments with the lowest expected health outcome. When the number of segments is large (e.g., 101 with discretization $p^i \in \{0\%,\ 1\%, \ldots, 100\%\}$), each segment typically uses only a small fraction of resources. In such cases, the maximin solution does not provide adequate guidance for allocating resources to segments other than the one that is worst off. In particular, the maximin problem has multiple optimal solutions of which most fail to utilize all available resources. To overcome this problem, we follow Luss (1999) and state that the allocation of possible leftover resources among patient segments $i \in \{1, \ldots, I\}$ is called *equitable*, if the expected health outcome of any

segment i cannot be improved without either violating a constraint or decreasing the expected health outcome of some other segment whose health outcome is either equal to or less than that of segment i.

An equitable solution to the maximin binary linear programming problem can be found by solving a lexicographic maximin problem (Luss, 1999). In particular, let $f(\boldsymbol{x}) = [f_{i_1}(\boldsymbol{x}_{i_1}), f_{i_2}(\boldsymbol{x}_{i_2}), \ldots, f_{i_I}(\boldsymbol{x}_{i_I})]$, where the elements $f_i(\boldsymbol{x}_i) = \sum_j x_{i,j_i} h_{i,j_i}$ corresponding to the expected health outcomes of different patient segments are arranged in a nondecreasing order such that $f_{i_1}(\boldsymbol{x}_{i_1}) \leq f_{i_2}(\boldsymbol{x}_{i_2}) \leq \ldots \leq f_{i_I}(\boldsymbol{x}_{i_I})$. The lexicographic maximin problem, referred to as PROBLEM E-ALLOCATION, is formulated as follows:

$$\begin{aligned} E^* = \quad & \text{lex max}_{\boldsymbol{x}}\, f(\boldsymbol{x}) \\ \text{subject to} \quad & \textstyle\sum_{i=1}^{I} \sum_{j_i=1}^{J_i} x_{i,j_i} d_i c_{i,j_i} \leq B \end{aligned} \tag{10}$$

$$\textstyle\sum_{j=1}^{J_i} x_{i,j_i} = 1 \text{ for all } i \in \{1, \ldots, I\} \tag{11}$$

$$x_{i,j_i} \in \{0, 1\} \text{ for all } i \in \{1, \ldots, I\},\ j_i \in \{1, \ldots, J_i\}. \tag{12}$$

The optimal solution is determined by repeatedly solving maximin binary linear programming problems with objective function $\max_{\boldsymbol{x}} \min_i x_{i,j_i} h_{i,j_i}$ and constraints (10)–(12). At each iteration, the optimal strategy is determined for some patient segment i_1 that is currently worst off. A new maximin problem is then formulated without this patient segment such that the level of available resources B in the right-hand side of (10) is decreased by the cost $d_{i_1} c_{i_1, j^*_{i_1}}$ of their optimal strategy. This process is continued until all resources have been exhausted. The resulting strategy vector $\boldsymbol{x}$ is equitable to all patient segments, in the sense that the expected health outcome cannot be increased for any segment without either violating a feasibility constraint or decreasing the expected health outcome for another segment that is already worse off (for more information, see Luss, 1999).

AN ILLUSTRATIVE EXAMPLE

Problem Description

We illustrate our model by revisiting the study on the optimal use of genetic testing in the prevention of CHD events (Hynninen, Linna, & Vilkkumaa, 2019). CHDs are the leading global cause of death, accounting for more than 7 million deaths per year (WHO, 2016). As primary prevention, a patient at risk of having a CHD event can be treated with statin medication. To target preventive interventions cost-effectively, it is important to obtain reliable prognostic information on the patient's state of health. Traditionally, risk measures such as the Framingham Risk Score (Anderson, Odell, Wilson, & Kannel, 1991; Wilson et al., 1998) or FINRISK function (Vartiainen et al., 2007), based on clinical factors and lipid measurements, have been used for this purpose. Over the last decade, much research has been done to assess the benefits of using genome information alongside traditional risk factors in the prevention of CHD.

Figure 4: A decision tree representing alternative testing strategies using traditional risk factors (TRS) and genetic testing (GRS).

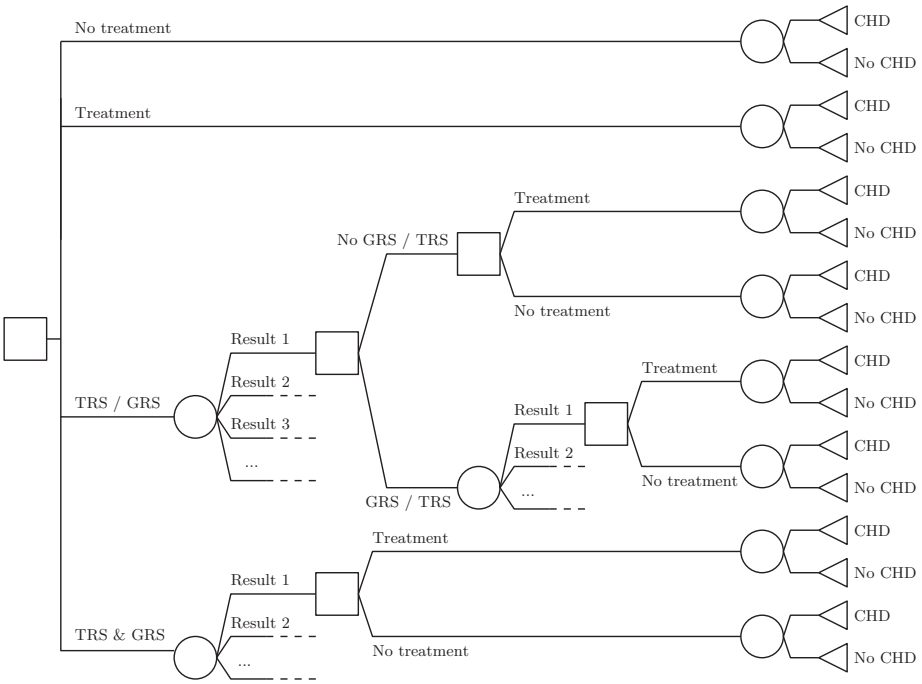

We optimize the use of traditional risk factors and genetic testing (either simultaneously or in subsequent stages) for obtaining prognostic information about whether a patient will ($S = 1$) or will not have ($S = 0$) a CHD event in the following 10 years. Based on the prognosis combined from the test results and prior information, the patient is either treated ($a = 1$) or not ($a = 0$) with statin medication. This process is illustrated by the decision tree in Figure 4.

In this article, test results on traditional risk factors are represented by Traditional Risk Score (TRS) that relies on information about the patient's gender, age, total cholesterol, high-density lipoprotein-cholesterol, systolic blood pressure, blood pressure treatment, smoking, prevalent diabetes, family history of myocardial infarction, and lipid treatment. The association between the factors of TRS and the rate of incident CHD events has been estimated by using a Cox proportional hazards model, which is adjusted for the traditional risk factors at baseline and uses age as the time scale. Hazard rate is assumed to remain constant, whereby the TRS is computed by using an exponential survival function. Similar risk scores have been used in Ripatti et al. (2010) and Tikkanen, Havulinna, Palotie, Salomaa, and Ripatti (2013). As the Genetic Risk Score (GRS) we use a novel score of 49,310 single-nucleotide polymorphisms (Abraham et al., 2016). Both TRS and GRS return a single number that can be interpreted as the probability of the patient having a CHD event in the following 10 years.

Table 1: Costs and health outcomes of CHD.

	Statin Treatment	No Statin Treatment
Costs of CHD (€)		
CHD event	12,058	14,629
No CHD event	1,927	0
Health outcomes (QALY)		
CHD event	7.143	6.952
No CHD event	7.689	7.706

Data

A CHD event is defined as (i) hospitalization caused by unstable angina (I200; ICD-10), acute myocardial infarction (I21), subsequent ST elevation and non-ST elevation myocardial infarction (I22), or revascularization event, or (ii) death caused by diagnosis I20-I25, I46, R96, R98 (ICD-10). Here, we present in brief the data on the costs and health outcomes of such events, the costs of tests and treatments, and the distribution of patients with different initial probabilities of having a CHD event in the following 10 years. More detailed descriptions of these data can be found in the Supporting Information and in Hynninen et al. (2019).

Table 1 shows the estimated costs and health outcomes for the occurrence of a CHD event. Here, the costs and health outcomes were measured from the perspective of the Finnish health care sector using a 10-year time horizon, and estimated based on national registers and literature reviews. The estimates were adjusted to 2015 level using the health care price index published by the Association of Finnish Local and Regional Authorities (2016). Following the U.S. Panel on Cost-Effectiveness in Health and Medicine, a 3% annual discount rate was applied to both health outcomes and costs (Gold et al., 1996).

The expected costs and health outcomes were first estimated for each year of the 10-year time horizon, and then discounted and summed together to represent the net present values of the costs and health outcomes shown in Table 1. Based on the FINRISK function (Vartiainen et al., 2007), the expected time of the CHD event was estimated to be 5.75 years given that the event would occur within the 10-year time horizon. The average annual cost of a CHD event was computed as a weighted average of the costs of fatal (22%) and nonfatal events (78%). The assumption of the mutual exclusiveness of fatal and nonfatal events was made to keep this illustrative model relatively simple, although in reality one or multiple nonfatal events could occur before a fatal event during the 10-year time horizon. The cost of a nonfatal CHD event was 19,860 €, which included the cost of treating an acute event in special health care (on average 9,015 €) and one-year follow-up in both primary and special health care (on average 10,844 €; National Discharge Register). We assumed that secondary prevention would be carried out for every survived patient beginning at the expected time of a nonfatal CHD event (5.75 years) until the end of the time horizon. The average annual cost of secondary prevention was estimated to be 451 € (Hujanen, Kapiainen, Tuominen, & Pekurinen, 2008). The cost of a fatal event was 2,417 €, which included the cost of only a few care

Table 2: Costs of testing.

	Cost (€)
Traditional Risk Score (TRS)	173
Genetic Risk Score (GRS)	200
TRS and GRS simultaneously	363

days (Hujanen et al., 2008; Aarnio, Korhonen, Huupponen, & Martikainen, 2015). The annual cost of statin treatment (226 €) was computed based on average annual medication cost (53 €) (Kiviniemi et al., 2011; Aarnio et al., 2015) and monitoring costs for one additional doctor, nurse, and laboratory visit annually (173 €), priced according to Finnish standard health care costs (Kapiainen, Väisänen, & Haula, 2014).

Health outcomes in Table 1 were estimated in QALY by applying the health-related quality-of-life (QoL) decrements (Kattainen, Sintonen, & Kettunen, 2005; Peura, Martikainen, Soini, Hallinen, & Niskanen, 2008; Koskinen, Lundqvist, & Ristiluoma, 2012) to each year spent in CHD-event states. Based on a meta-analysis (Taylor, Huffman, & Ebrahim, 2013), statins were estimated to reduce the risk for a fatal or nonfatal CHD event by 27%. For each year of statin therapy, a small QoL decrement of 0.002 was applied (Gage, Cardinalli, & Owens, 1996; Hutchins, Viera, Sheridan, & Pignone, 2015).

Table 2 shows the testing and treatment costs. The cost of carrying out TRS was estimated to be 173 €, including doctor and nurse visits and a blood panel, based on Finnish standard health care costs (Kapiainen et al., 2014). The cost of GRS was assumed to be 200 € based on discussions with subject experts. If TRS and GRS were to be carried out simultaneously, the combined cost of these tests was assumed to be 363 €, reflecting a cost saving of 10 € from needing to take the blood or tissue sample only once instead of twice (i.e., the cost of an additional testing stage $c^{\text{stage}} = 10$; the cost is illustrative and reflects the price level in Finland at the time).

The data on the accuracy of testing were derived from FINRISK studies with a total of 17,457 subjects (1992, 1997, 2002, and 2007 cohorts; Borodulin et al., 2015) with known TRS and GRS results, as well as the observed 919 CHD events during the 10-year follow-up period. We assumed that TRS and GRS would be conditionally independent, meaning that the result of each test would depend only on the patient's state of health and not on the result of the other test. This assumption was justified by the negligible correlation between TRS factors and GRS in the data provided by the Finnish Institute for Molecular Medicine: the largest absolute value of this correlation was .09.

Finally, Figure 5 shows the initial disease probability distribution d_i (i.e., the distribution of the patients over different initial probabilities p^i of having a CHD event in the following 10 years) among 100,000 nonsmoking Finnish men and women aged 45 years or older who had no diabetes or family history of CHDs. This distribution was based on the FINRISK function, which is a gender-specific logistic regression function on risk factors such as age, smoking, systolic blood pressure, and HDL-cholesterol (Vartiainen et al., 2007). For illustrative purposes, Figure 5

Figure 5: Risk distribution of 100,000 Finnish people aged 45 years or older.

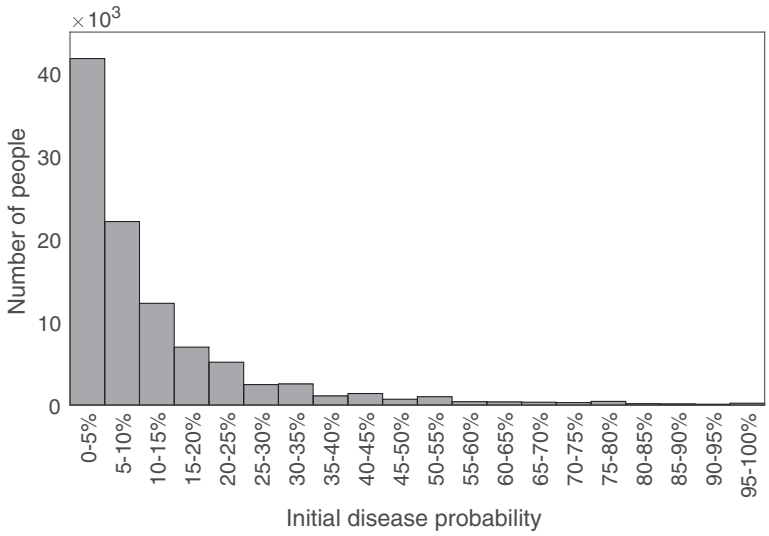

presents the histogram for aggregated initial disease probabilities $p^i \in \{0-5\%, 5-10\%,...,95-100\%\}$. The numerical results in this case study were, nevertheless, computed with a denser discretization $p^i \in \{0\%,\ 1\%, \ldots, 100\%\}$.

Optimal Population-Level Strategies

The optimal strategies were solved using MATLAB 2017a on a standard laptop (1.9 GHz, 8 GB memory). First, the sets of Pareto optimal strategies were identified for each initial disease probability p_0. We used 101 discretization points for p_0 $(0\%, 1\%, \ldots, 99\%, 100\%)$ and $J = 500$ discretization points between minimum ($b_1 = 0$€) and maximum ($b_J = 12,100$€) values for the total costs corresponding to a single patient. This number of discretization points provided detailed results with reasonable computation time (in total 418 minutes). The maximum value b_J was based on the maximum cost of 12,058 € per patient (i.e., the cost of treating a patient with the disease, see Table 1) identified in initial test runs with a lower number of discretization points. Second, population-level strategies were optimized for the utilitarian (PROBLEM U-ALLOCATION) and egalitarian objective (PROBLEM E-ALLOCATION) for 49 values of total cost level $B \in \{160, 161, \ldots, 208\}$M€, where B includes all testing, treatment, and posttreatment costs. The lowest value $B = 160$M€ corresponds to the lowest possible expected cost for the health care system, and the highest value $B = 208$M€ to the expected cost of applying the strategy with maximum expected health outcome to each individual patient. The total computation time of this second phase was 31 seconds for the utilitarian approach and 87 seconds for the egalitarian approach.

Figures 6(a) and (b) show the optimal testing strategies and the corresponding shares of the population for utilitarian and egalitarian objectives at different levels of total cost B. At the lowest and highest total cost levels, the optimal strategies

Figure 6: Optimal testing strategies in various total cost levels.

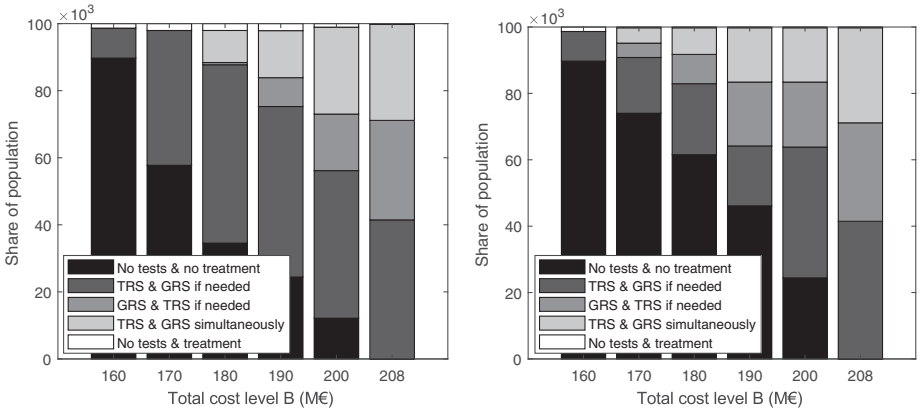

(a) Utilitarian strategy (b) Egalitarian strategy

would be the same for both population-level objectives in these figures and all the later ones as well, because either the least expensive or the most expensive strategy would be chosen for each segment, respectively. At the highest cost level, everyone in the population would be tested. Nevertheless, even at the lowest cost level it would be optimal to carry out TRS to 9% of the population and, if the results remain inconclusive, also GRS. This reflects the need to avoid the costs of not treating patients with a moderately low initial disease probability who are, in fact, at risk of having a CHD event in the following 10 years. Between these extreme cost levels, the utilitarian strategies would subject many more patients to tests, especially to TRS. However, a small segment of patients with high initial probability of disease would be treated without testing. In contrast, the egalitarian strategies would focus on finding those patients in the smaller segments with high initial disease probabilities who in fact are not at risk of having a CHD event in the next 10 years. To prevent the unnecessary treatment of these patients, efforts would be taken to obtain accurate information about their risk through the use of both TRS and GRS. In particular, the use of the more expensive GRS (either before TRS or simultaneously with it) would be optimal for a notable share of the population even at relatively low cost levels.

Figures 7(a) and (b) show the optimal prioritization of tests for patient segments in the utilitarian and egalitarian approaches, respectively. Here, the color of a cell represents the expected testing cost for a single patient in a given segment defined by the 5% interval of initial disease probability (vertical axis) at a given cost level B (horizontal axis). A white cell indicates that the patient segment would not be tested at the selected cost level. On the other hand, a black cell indicates the patient segment would be tested maximally so that both TRS and GRS would be carried out to all patients in it. If the cell is gray, then some but not all tests would be carried out within the patient segment at the specific cost level. When interpreting these figures, it should be noted that most patients belong to segments with low

Figure 7: Optimal testing in patient segments. The color shows the average amount of resources allocated to a single patient with a particular initial disease probability at a given cost level.

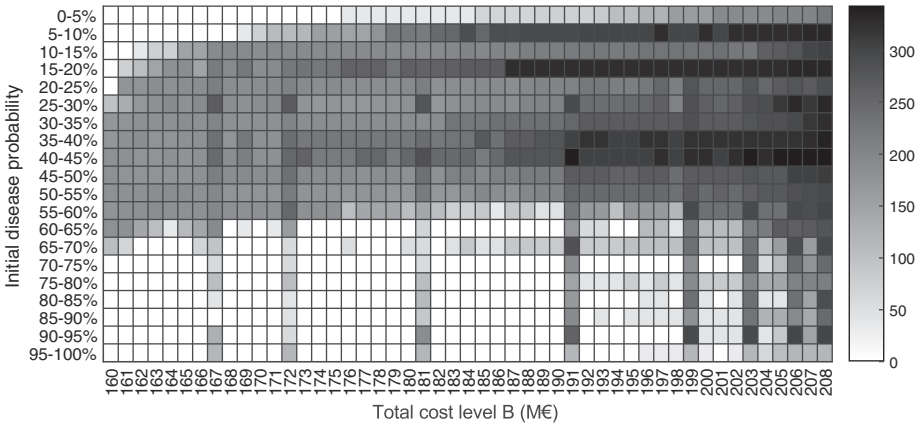

(a) Utilitarian objective

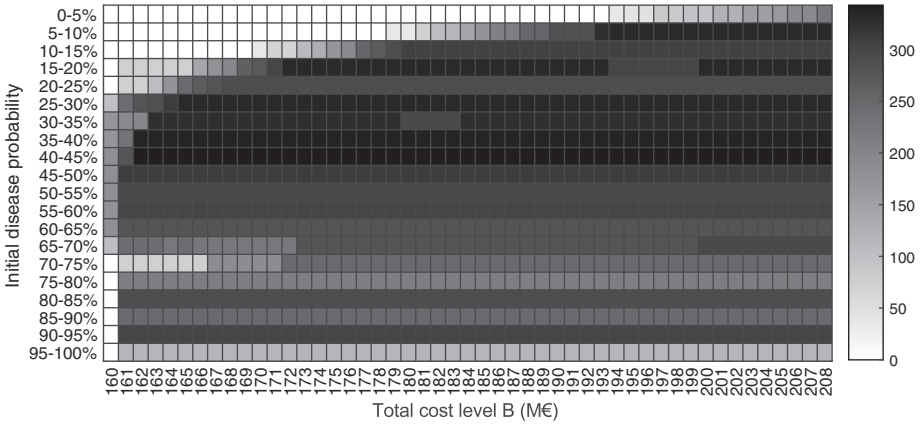

(b) Egalitarian objective

initial disease probabilities: 83% of the patients have an initial disease probability of 20% or lower.

The utilitarian approach (Figure 7a) would not test patient segments with low and high initial disease probabilities so that treatment decisions for these segments would be made solely based on the initial disease probabilities. This means that tests would be focused on segments with intermediate initial disease probabilities. This is in line with Delquié's (2008) result that states that the value of information is highest when the decision-maker is initially indifferent between alternatives (here, treating vs. not treating). With more resources, it would become possible to test more patient segments with lower initial disease probabilities. In the egalitarian approach (Figure 7b), in contrast, tests would be focused on patients with a high

Figure 8: Optimal utilitarian strategy at cost level 169M€.

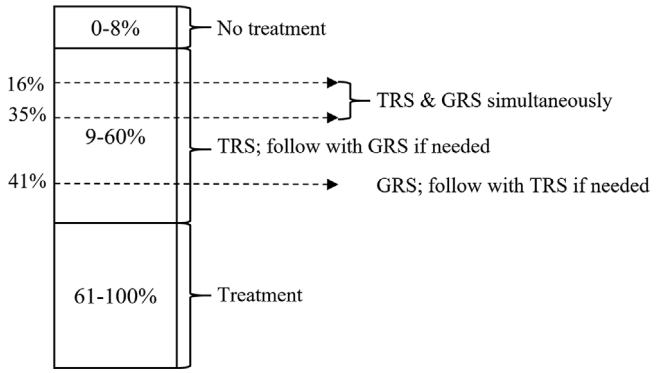

or intermediate initial disease probability. This strategy reflects the egalitarian objective of prioritizing the maximization of the expected health outcome for those patients whose expected health outcome is the lowest in that it seeks to avoid the unnecessary treatment of patients with high initial disease probabilities who are, in fact, not at risk of having a CHD event in the following 10 years. On the other hand, for the vast majority of the population with low initial disease probability (i.e., a high expected health outcome), the decision of not to treat would be made without tests, unless the total cost level B was sufficiently high.

At each cost level, trade-offs would be made between patient segments, tests, and treatments in order to maximize the stated objective. These trade-offs explain the horizontal deviations (i.e., deviations between cost levels within a given patient segment) that can be seen in Figure 7(a) at, for example, cost levels 167M€, 172M€, 181M€, and 191M€. At these cost levels, the increase in available resources would not yet enable the most efficient option to improve the expected health outcome that, in these cases, would be to carry out more tests for some relatively large patient segment. Instead, additional resources would be used to cover the more extensive testing of smaller patient segments with high initial disease probabilities.

The vertical deviations in Figures 7(a) and (b) (i.e., deviations between patient segments within a given cost level) are related to using real, nonsmoothed data to estimate the accuracy of testing as well as the population distribution. This resulted in situations where (i) one test or test combination would be particularly beneficial for some small patient segment but not for the adjacent segments, and (ii) the sizes of patient segments would vary substantially, especially for large disease probabilities. Figure 8, for example, shows the utilitarian optimal strategies at cost level 169M€. These strategies are nonconvex in that, for example, a strategy that would be optimal for initial probabilities 15% and 17% (namely, to carry out TRS and then GRS only if needed) would not be optimal for initial probability 16%, for which it would pay off to carry out TRS and GRS simultaneously. Such discontinuities in optimal strategies highlight the need to compute these strategies for each initial disease probability separately.

Figure 9: Optimal prioritization of treatments for patient segments. The color shows the average amount of resources for a single patient with a given initial disease probability.

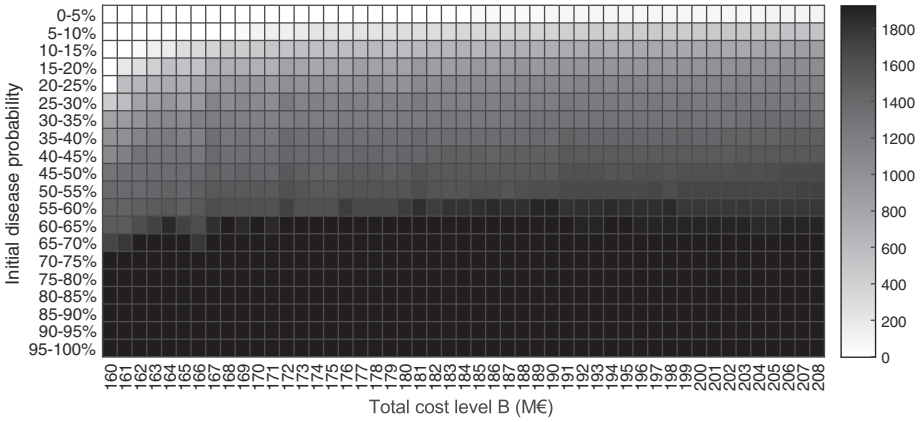

(a) Utilitarian objective

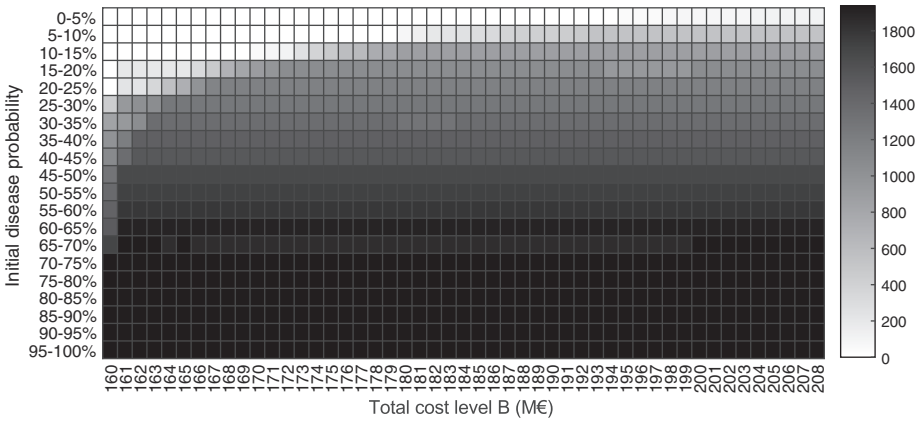

(b) Egalitarian objective

Figures 9(a) and (b) show the optimal prioritization of treatment resources between patients under utilitarian and egalitarian objectives, respectively. The interpretation is similar to Figures 7(a) and (b): a white cell implies that no patient in the particular patient segment would be treated at the given cost level, a black cell implies that all patients in the segment would be treated, and a gray cells implies that some patients in the segment would be treated but not all. The main difference in the allocation of treatment resources between different approaches is that at low cost levels B, the egalitarian approach would allocate less resources to treating patients with low initial disease probability. This is because these patients would not be tested and the decision to not treat would be made based on initial disease probability alone (cf. Figure 7b).

Figure 10: The expected health outcome of an individual patient in various total cost levels.

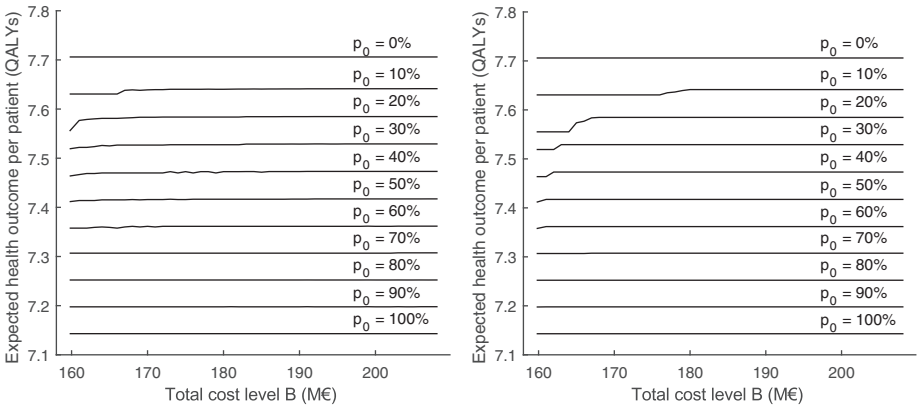

(a) Utilitarian strategy

(b) Egalitarian strategy

Figures 10(a) and (b) show how the expected health outcome per patient changes in the utilitarian and egalitarian approaches, respectively, as a function of the total cost level. In both figures, the expected health outcomes are shown for patients with initial probabilities $p_0 = 0, 10\%, \ldots, 90\%, 100\%$. In the utilitarian approach (Figure 10a), additional resources would be allocated to patient segments in which they would yield the highest increase in the expected health outcome. For instance, at low initial cost levels, additional resources would be given to the segments defined by intermediate probabilities $p_0 = 20\%$, 30%, and 40% that are close to the treatment threshold. The egalitarian approach (Figure 10b) would assign additional resources to a patient segment only if the expected health outcomes of other segments with higher initial disease probabilities had already been raised to their maximum levels.

Figure 11 shows the cumulative consumption of resources by patients in the population at four different cost levels (160, 175, 190, and 208M€) for the utilitarian and egalitarian objectives. At the lowest cost level $B = 160$M€, the utilitarian and egalitarian strategies are the same, and all resources would be consumed by approximately 12% of the population. At intermediate cost levels $B = 175$M€ and $B = 190$M€, the utilitarian approach would distribute resources among more patients than the egalitarian one. For instance, when $B = 175$M€, 36% of patients would use all resources in the egalitarian approach, whereas in the utilitarian approach these resources would be consumed by 58% of patients. The differences between the utilitarian and egalitarian approach at intermediate cost levels are indicated by the gray areas. At cost level $B = 208$M€, all patients could be tested and treated to the point of maximizing their expected health outcome, and the utilitarian and egalitarian approaches would coincide.

Finally, Figure 12 shows the expected population-level health outcome H as a function of the expected population-level cost B for the utilitarian and egalitarian approaches (cf. production possibility frontier; Hutubessy et al., 2003). This figure helps compare the results of our model to those of traditional CEA. For instance,

Figure 11: The distribution of total costs on various shares of population.

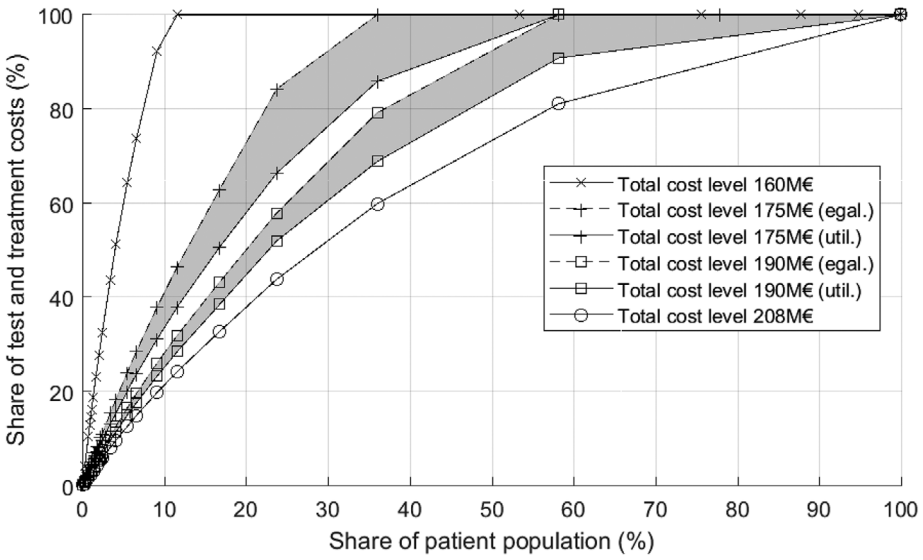

Figure 12: Health outcome to costs of optimal utilitarian and egalitarian strategies at different cost levels for a population of 100,000 people.

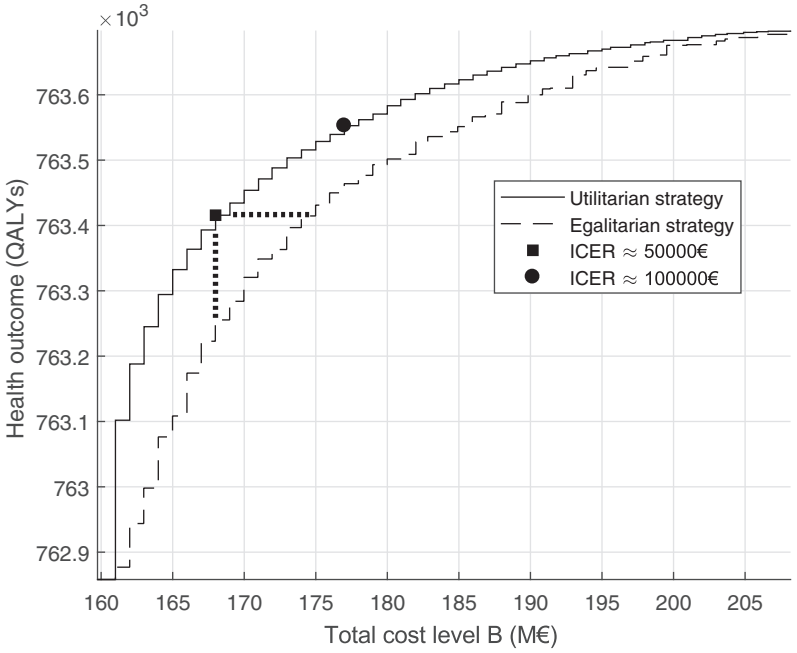

the central evaluation measure in CEA is the incremental cost-effectiveness ratio (ICER), which describes the marginal cost of an additional quality-adjusted life year (Gold et al., 1996). Based on Figure 12, the ICER can be computed at all cost levels B as the inverse of the slope of the curve: $1/H'(B)$. The lower the ICER, the more cost-effective the strategy. For the utilitarian objective, the curve $H(B)$ has a decreasing slope, whereby the ICER is increasing in the total cost level B. In other words, the higher the total cost level, the less cost-effective it would be to spend an additional unit of resources.

In Finland, no explicit ICER threshold values are commonly used to support cost-effective health care intervention investments. Elsewhere, values such as \$50,000 and \$100,000 per QALY gained (in the United States) or £20000 and £30000 per QALY gained (in the UK) have been used (Cleemput, Neyt, Thiry, De Laet, & Leys, 2008; Neumann, Cohen, & Weinstein, 2014). As an illustrative example, Figure 12 shows the total cost levels B at which the ICER for the utilitarian strategy would be approximately 50,000 €/QALY (black square) and 100,000 €/QALY (black circle). If, for instance, the threshold value of 50,000 €/QALY was used, strategies with total cost levels above $B = 168$M€ would not be cost-effective and should not be considered by a rational decision maker. For egalitarian strategies, the ICER does not increase monotonically, whereby similar conclusions cannot be made.

Figure 12 shows that at all cost levels B, the egalitarian strategy would give a lower expected health outcome than the utilitarian strategy. This is because the objective in the egalitarian strategy is not to maximize the total expected health outcome, but to maximize the well-being of those who are worst off. Hence, resources might be allocated to expensive interventions which would increase the health outcome only marginally compared to their costs. Comparing the expected costs and health outcomes of utilitarian and egalitarian strategies would help policy makers assess the cost of equity, defined as the difference between (i) expected health outcomes at a given cost level or (ii) cost levels required for a given expected health outcome (e.g., Stinnett & Paltiel, 1996; Weatherly et al., 2009). For example, compared to the optimal utilitarian strategy corresponding to ICER $\approx 50{,}000$€, the cost of equity would be roughly (i) 160 QALYs to maintain the cost level $B = 168$M€, or (ii) 6M€ to maintain the expected health outcome of 763,420 QALYs (dotted lines in Figure 12).

DISCUSSION AND CONCLUSIONS

We have developed a decision model to optimize the test and treatment strategies for different patient segments, subject to a population-level resource constraint and in view of two population-level objectives: maximizing the expected health outcome of the population (utilitarian) and maximizing the expected health outcome of those who are worst off (egalitarian). Among other uses, our model can be useful in preparing for emergencies and allocating resources to humanitarian health care, where the primary goal is to improve the population-level health rather than that of any particular individual. In such settings, the model can be used to generate defensible policy recommendations for (i) segmenting and prioritizing

patients, (ii) anticipating how much resources are needed for tests and treatments, and (iii) assessing the cost-effectiveness of new tests and treatments.

The model has several advantages over earlier approaches to the evaluation of health care programs and intervention strategies. First, instead of assuming that the population is homogenous or consists of few large and predetermined subpopulations only, the optimal testing and treatment strategies are computed explicitly for all risk levels associated with different initial probabilities of having a given disease. With multiple tests, test results, and testing stages, this means that the results can be used for the specification of meaningful patient segments within which similar treatment and testing strategies are applied. Furthermore, because the use of new tests and treatment options can be optimized along with existing ones, the model helps assess if investments into these options are cost-effective and, if so, how they should be used.

Second, unlike earlier approaches, the model guides the interpretation of imperfect test results and the ensuing selection of treatments for each patient segment and resource level. This can be particularly helpful with nonbinary test results that may not provide immediate suggestions as to which treatments should be selected. That is, the model helps interpret test results so that this interpretation, too, contributes to the chosen population-level objective.

Third, the model helps ensure that the health outcomes are indeed maximized subject to relevant resource constraints, because the multi-objective optimization identifies Pareto optimal strategies for which no other strategy would offer a higher expected health outcome at a lower expected cost. These optimal strategies are determined without monetizing health outcomes by employing parameters such as the highly contested WTP threshold (e.g., Neumann et al., 2014). Finally, contrasting the utilitarian and egalitarian objectives helps explore how the optimal testing and treatment strategies and corresponding resource allocations to patient segments depend on policy objectives. For instance, in our case study on the prevention of CHD events, the egalitarian approach suggested spending more resources on obtaining prognostic information about the few patients who have a high initial probability of CHD, while the recommendation for many patients with low initial probability was to do nothing. Although such a strategy would help avoid unnecessary treatments of patients who are worst off, it could lead to expensive interventions with marginal health benefits. More generally, our model permits comparisons between the costs and health outcomes of utilitarian and egalitarian policies and thus helps understand the cost of equity.

A key implication of our results is that policy-level objectives and resource constraints need to be considered jointly in developing recommendations for the allocation of resources between patient segments and the operational use of these resources for interventions such as tests and treatments within these segments. In our analysis, we have purposely contrasted utilitarian and egalitarian objectives to demonstrate differences in their implications. In practice, they could be balanced by excluding egalitarian strategies with an excessively high ICER (e.g., higher than 100,000 €/QALY) before optimizing population-level strategies in phase 2. This would prioritize health outcomes for those who are worst off, but not at unreasonably high costs. Alternatively, one could require that the objective function of the egalitarian approach must exceed a given threshold level and determine

cost-effective strategies by maximizing the utilitarian objective function subject to such a constraint (Hooker & Williams, 2012). A further approach would be to build a multiattribute model in which the trade-off between utilitarianism and egalitarianism is treated with attribute weights.

Another important takeaway is that the optimal resource allocation—which is operationalized by determining optimal testing and treatment strategies for different patient segments—depends on (i) the distribution of patients in segments representing different risk levels, (ii) the choice of the population-level objective (i.e., utilitarian vs. egalitarian), and (iii) the available resources. Indeed, it is only by considering these three aspects jointly that one can reach conclusive statements about the cost-effectiveness of tests, for instance. In preparing for health care emergencies, information about the risk distribution may not be available; but alternative assumptions about this distribution could be formulated as scenarios for assessing what resources would be needed to reach the chosen population-level objectives satisfactorily. The full-scale model could also be used as benchmark for assessing more straightforward resource allocations (i.e., fewer patient segments, tests, or testing stages) that may be more viable in situations with considerable time pressure.

There are technical assumptions of our model that could be relaxed. First, the calculation of posterior probabilities was premised on the assumption that the results of different tests are conditionally independent, given the patient's state of health. If this is not the case, the model will not give fully accurate results (van Walraven, Austin, Jennings, & Forster, 2009; Novielli, Cooper, & Sutton, 2013), but these results could still provide a viable starting point if there is no information about dependencies (Gelman, Carlin, Stern, & Rubin, 2014). Moreover, optimal strategies for conditionally independent tests can be explored to gain insights into which tests could be optimal in the presence of dependencies, allowing efforts on studying dependencies to be focused accordingly.

Second, the presented decision model is static in that it does not account for the possible deterioration of health. Technically, the evolution of a patient's state of health can be modeled, for instance, with time-dependent state transition models with the aim of guiding time-dependent resource allocation decisions. A third, more challenging extension is that of accommodating multimorbidity, recognizing that many diseases (such as type II diabetes and coronary heart diseases) are interdependent and can often be prevented by similar interventions (Eranti et al., 2016). Although multimorbidity could, in principle, be modeled by treating the patient's state as a vector-valued variable, the resulting data requirements would be substantially higher. In theory, a model in which the state would capture the overall health conditions of each member of the population could even support the allocation of total health care resources. Yet, in the absence of the required data, we have decided to limit the scope of this article to a single disease.

A notable limitation of our model is that patients are segmented solely based on their probability of developing a disease. Yet in reality, two patients with the same probability of disease may respond to tests and treatments differently if they are different in terms of characteristics such as smoking or obesity. From the point of health equity, groupings based on race, gender, or socio-economic status may also be of interest. Technically, it is straightforward to extend our model to account for any pregrouping of patients based on such characteristics. That is, the nondom-

inated strategies in phase 1 can be identified separately for each pregroup, based on probability of disease. Then, the population-level resources in phase 2 can be allocated between the segments defined by both their initial disease probability and the characteristics employed in pregrouping. In practice, such pregrouping of patients would require that all model parameters (including health outcomes and costs as well as the probabilities of different test results) are estimated separately for each pregroup. As a result, the data requirements would increase exponentially as a function of how many characteristics are considered. Although such data are not yet readily available, there are notable efforts to collect and analyze large individual-level data sets to support the development of patient-specific intervention strategies (cf. *precision medicine*; National Research Council, 2011). In Finland, for instance, the large-scale FinnGen study will collect 500,000 blood samples, combining genome information with digital health care data from national health registries (FinnGen, 2018). The government will also establish a genome centre that will administer a national database of genomes and promote the use of genetic information in health care, research, and innovation activities (Ministry of Social Affairs and Health (Finland), 2017). Meanwhile, expert panels could be consulted to obtain relevant estimates for the cost-effective targeting of health care resources.

Finally, our model is computationally demanding. The computation time depends primarily on the number of tests, the number of possible test results, and the number of points in resource level discretization. Because the problem of finding optimal test and treatment strategies is nonconvex, it is difficult to establish structural properties for speeding up computations. Instead, we used the ε-constraint method to approximate the set of Pareto optimal strategies. The computation time could be reduced by using more efficient methods, such as heuristics or evolutionary algorithms (e.g., Shukla & Deb, 2007; Rauner, Gutjahr, Heidenberger, Wagner, & Pasia, 2010). However, because the purpose of the model is to support macro-level decision making (e.g., developing clinical care guidelines, assessing testing or treatment technologies), results need not be recomputed often. Nevertheless, improved computational algorithms would be helpful, especially for the implementation of the model extensions outlined above.

ACKNOWLEDGMENTS

We thank the anonymous reviewers and the editors for their comments which improved the quality, readability, and relevance of this article. We also thank Miika Linna for insightful comments and expertise. This research has been supported by the Foundation for Economic Education (Grant no. 14-7589) (Vilkkumaa) and the Research project Platform Value Now, funded by the Grant 314207 of the Strategic Research Council of the Academy of Finland (Salo).

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Supporting Information

References

Aarnio, E., Korhonen, M. J., Huupponen, R., & Martikainen, J. (2015). Cost-effectiveness of statin treatment for primary prevention in conditions of real-world adherence – Estimates from the Finnish prescription register. *Atherosclerosis*, *239*(1), 240–247.

Abraham, G., Havulinna, A. S., Bhalala, O. G., Byars, S. G., De Livera, A. M., Yetukuri, L., et al. (2016). Genomic prediction of coronary heart disease. *European Heart Journal*, *37*(43), 3267–3278.

Anderson, K. M., Odell, P. M., Wilson, P. W., & Kannel, W. B. (1991). Cardiovascular disease risk profiles. *American Heart Journal*, *121*(1), 293–298.

Association of Finnish Local and Regional Authorities. (2016). Health care price index of Finnish local and regional authorities.

Ayer, T., Alagoz, O., & Stout, N. K. (2012). Or forum - a POMDP approach to personalize mammography screening decisions. *Operations Research*, *60*(5), 1019–1034.

Beauchamp, T., & Childress, J. (2001). *Principles of biomedical ethics* (5th ed.). New York, NY: Oxford University Press.

Bellman, R. E. (1957). *Dynamic programming*. Princeton, NJ: Princeton University Press.

Ben-Assuli, O., & Leshno, M. (2013). Using electronic medical records in admission decisions: A cost effectiveness analysis. *Decision Sciences*, *44*(3), 463–481.

Blanchet, K., Sistenich, V., Ramesh, A., Frison, S., Warren, E., Hossain, M., et al. (2013). *An evidence review of research on health interventions in humanitarian crises*. London, UK: London School of Hygiene & Tropical Medicine.

Borodulin, K., Vartiainen, E., Peltonen, M., Jousilahti, P., Juolevi, A., Laatikainen, T., et al. (2015). Forty-year trends in cardiovascular risk factors in Finland. *The European Journal of Public Health*, *25*(3), 539–546.

Brandeau, M. L., Zaric, G. S., & Richter, A. (2003). Resource allocation for control of infectious diseases in multiple independent populations: Beyond cost-effectiveness analysis. *Journal of Health Economics*, *22*(4), 575–598.

CADTH (2017). *Guidelines for the economic evaluation of health technologies: Canada* (4th ed.). CADTH.

Cantor, S. B. (1994). Cost-effectiveness analysis, extended dominance, and ethics: A quantitative assessment. *Medical Decision Making*, *14*(3), 259–265.

Cevik, M., Ayer, T., Alagoz, O., & Sprague, B. L. (2018). Analysis of mammography screening policies under resource constraints. *Production and Operations Management*, *27*(5), 949–972.

Chalabi, Z., Epstein, D., McKenna, C., & Claxton, K. (2008). Uncertainty and value of information when allocating resources within and between healthcare programmes. *European Journal of Operational Research*, *191*(2), 530–539.

Cleary, S., Mooney, G., & McIntyre, D. (2010). Equity and efficiency in HIV-treatment in South Africa: The contribution of mathematical programming to priority setting. *Health Economics*, *19*(10), 1166–1180.

Cleemput, I., Neyt, M., Thiry, N., De Laet, C., & Leys, M. (2008). *Threshold values for cost-effectiveness in health care*. Health Technology Assessment (HTA). Belgian Health Care Knowledge Centre Reports 100.

Cuadras-Morató, X., Pinto-Prades, J.-L., & Abellán-Perpiñán, J.-M. (2001). Equity considerations in health care: The relevance of claims. *Health Economics*, *10*(3), 187–205.

Das, I., & Dennis, J. E. (1997). A closer look at drawbacks of minimizing weighted sums of objectives for Pareto set generation in multicriteria optimization problems. *Structural Optimization*, *14*(1), 63–69.

Deb, K. (2014). Multi-objective optimization. In Burke, E. K. & Kendall, G., (Eds), *Search methodologies - introductory tutorials in optimization and decision support techniques*. New York, NY: Springer Science & Business Media, 403–449.

Delquié, P. (2008). The value of information and intensity of preference. *Decision Analysis*, *5*(3), 129–139.

Demarteau, N., Breuer, T., & Standaert, B. (2012). Selecting a mix of prevention strategies against cervical cancer for maximum efficiency with an optimization program. *PharmacoEconomics*, *30*(4), 337–353.

Deo, S., Rajaram, K., Rath, S., Karmarkar, U. S., & Goetz, M. B. (2015). Planning for HIV screening, testing, and care at the veterans health administration. *Operations Research*, *63*(2), 287–304.

Dolan, P., & Cookson, R. (2000). A qualitative study of the extent to which health gain matters when choosing between groups of patients. *Health Policy*, *51*(1), 19–30.

Drummond, M. F., Sculpher, M. J., Torrance, G. W., O'Brien, B. J., & Stoddart, G. L. (2005). *Methods for the economic evaluation of health care programmes* (3rd ed.). New York, NY: Oxford University Press.

Earnshaw, S. R., Richter, A., Sorensen, S. W., Hoerger, T. J., Hicks, K. A., Engelgau, M., et al. (2002). Optimal allocation of resources across four interventions for type 2 diabetes. *Medical Decision Making*, *22*(Suppl), S80–S91.

Eddy, D. M. (1991). Clinical decision making: from theory to practice. The individual vs society. Resolving the conflict. *JAMA*, *265*(18), 2396–2405.

Edejer, T. T.-T., Baltussen, R., Adam, T., Hutubessy, R., Acharya, A., Evans, D., et al. (2003). *Making choices in health: WHO guide to cost-effectiveness analysis*. Geneva, Switzerland: World Health Organization.

Emanuel, E. J. (2000). Justice and managed care: Four principles for the just allocation of health care resources. *The Hastings Center Report*, *30*(3), 8–16.

Epstein, D. M., Chalabi, Z., Claxton, K., & Sculpher, M. (2007). Efficiency, equity, and budgetary policies: Informing decisions using mathematical programming. *Medical Decision Making*, *27*(2), 128–37.

Eranti, A., Kerola, T., Aro, A. L., Tikkanen, J. T., Rissanen, H. A., Anttonen, O., et al. (2016). Diabetes, glucose tolerance, and the risk of sudden cardiac death. *BMC Cardiovascular Disorders*, *16*(1), 51.

FinnGen (2018). FinnGen research project.

Gage, B. F., Cardinalli, A. B., & Owens, D. K. (1996). The effect of stroke and stroke prophylaxis with aspirin or warfarin on quality of life. *Archives of Internal Medicine*, *156*(16), 1829–1836.

Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2014). *Bayesian data analysis* (2nd ed.). Boca Raton, FL: Chapman & Hall/CRC.

Gold, M., Siegel, J., Russell, L., & Weinstein, M. (Eds.) (1996). *Cost-effectiveness in health and medicine*. New York, NY: Oxford University Press.

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

Hooker, J. N., & Williams, H. P. (2012). Combining equity and utilitarianism in a mathematical programming model. *Management Science*, *58*(9), 1682–1693.

Hujanen, T., Kapiainen, S., Tuominen, U., & Pekurinen, M. (2008). *Terveydenhuollon yksikkökustannukset Suomessa vuonna 2006*. Stakes.

Hunink, M., Glasziou, P., Siegel, J. E., Weeks, J. C., Pliskin, J. S., Elstein, A. S., et al. (2001). *Decision making in health and medicine: Integrating evidence and values*. Cambridge, UK: Cambridge University Press.

Hutchins, R., Viera, A. J., Sheridan, S. L., & Pignone, M. P. (2015). Quantifying the utility of taking pills for cardiovascular prevention. *Circulation: Cardiovascular Quality and Outcomes*, *8*(2), 155–163.

Hutubessy, R., Chisholm, D., & Edejer, T. T. (2003). Generalized cost-effectiveness analysis for national-level priority-setting in the health sector. *Cost Effectiveness and Resource Allocation*, *1*(1), 8.

Hynninen, Y., Linna, M., & Vilkkumaa, E. (2019). Value of genetic testing in the prevention of coronary heart disease events. *PLoS ONE*, *14*(1), e0210010.

Johannesson, M., & Weinstein, M. C. (1993). On the decision rules of cost-effectiveness analysis. *Journal of Health Economics*, *12*(4), 459–467.

Johri, M., & Norheim, O. F. (2012). Can cost-effectiveness analysis integrate concerns for equity? Systematic review. *International Journal of Technology Assessment in Health Care*, *28*(02), 125–132.

Kapiainen, S., Väisänen, A., & Haula, T. (2014). *Terveyden- ja sosiaalihuollon yksikkökustannukset Suomessa vuonna 2011*. Number 3. THL.

Kattainen, E., Sintonen, H., & Kettunen, R. (2005). Health-related quality of life of coronary artery bypass grafting and percutaneous transluminal coronary artery angioplasty patients: 1-year follow-up. *International Journal of Technology Assessment in Health Care*, *2*(21), 172–179.

Kiviniemi, V., Peura, P., Helin-Salmivaara, A., Martikainen, J. E., Hartikainen, J., Huupponen, R., et al. (2011). Suboptimal use of statins at treatment initiation. *European Journal of Clinical Pharmacology*, *67*(9), 971–973.

Koskinen, S., Lundqvist, A., & Ristiluoma, N. (2012). *Terveys, toimintakyky ja hyvinvointi Suomessa 2011*. THL.

Laumanns, M., Thiele, L., & Zitzler, E. (2006). An efficient, adaptive parameter variation scheme for metaheuristics based on the epsilon-constraint method. *European Journal of Operational Research*, *169*(3), 932–942.

Lee, E., Lavieri, M. S., & Volk, M. (2019). Optimal screening for hepatocellular carcinoma: A restless bandit model. *Manufacturing & Service Operations Management*, *21*(1), 198–212.

Luss, H. (1999). On equitable resource allocation problems: A lexicographic minimax approach. *Operations Research*, *47*(3), 361–378.

Ministry of Social Affairs and Health (Finland). (2017). Establishment of a genome centre proceeds – the working group's memorandum sent for comments.

Murray, C. J., Evans, D. B., Acharya, A., & Baltussen, R. (2000). Development of WHO guidelines on generalized cost-effectiveness analysis. *Health Economics*, *9*, 235–251.

National Institute for Health and Care Excellence. (2018). Developing NICE guidelines: Incorporating economic evaluation.

National Research Council. (2011). *Toward precision medicine: Building a knowledge network for biomedical research and a new taxonomy of disease*. Washington, DC: National Academies Press.

Neumann, P. J., Cohen, J. T., & Weinstein, M. C. (2014). Updating cost-effectiveness – the curious resilience of the $50,000-per-QALY threshold. *New England Journal of Medicine*, *371*(9), 796–797.

Nord, E., Richardson, J., Street, A., Kuhse, H., & Singer, P. (1995). Maximizing health benefits vs egalitarianism: An Australian survey of health issues. *Social Science & Medicine*, *41*(10), 1429–1437.

Novielli, N., Cooper, N. J., & Sutton, A. J. (2013). Evaluating the cost-effectiveness of diagnostic tests in combination: is it important to allow for performance dependency?*Value in Health*, *16*(4), 536–541.

Peura, P., Martikainen, J., Soini, E., Hallinen, T., & Niskanen, L. (2008). Cost-effectiveness of statins in the prevention of coronary heart disease events in middle-aged Finnish men. *Current Medical Research and Opinion*, *24*(6), 1823–1832.

Rauner, M. S., Gutjahr, W. J., Heidenberger, K., Wagner, J., & Pasia, J. (2010). Dynamic policy modeling for chronic diseases: metaheuristic-based identification of pareto-optimal screening strategies. *Operations Research*, *58*(5), 1269–1286.

Rawls, J. (1971). *A theory of justice*. Cambridge, MA: Harvard University Press.

Ripatti, S., Tikkanen, E., Orho-Melander, M., Havulinna, A. S., Silander, K., Sharma, A., et al. (2010). A multilocus genetic risk score for coronary heart disease: Case-control and prospective cohort analyses. *The Lancet*, *376*(9750), 1393–1400.

Sabik, L. M., & Lie, R. K. (2008). Priority setting in health care: lessons from the experiences of eight countries. *International Journal for Equity in Health*, *7*(1), 4.

Sassi, F., Archard, L., & Le Grand, J. (2001). Equity and the economic evaluation of healthcare. *Health Technology Assessment*, *5*(3), 1–138.

Severens, J., Sonke, G., Laheij, R., Verbeek, A., & De Vries Robbé, P. (2001). Efficient diagnostic test sequence: Applications of the probability-modifying plot. *Journal of Clinical Epidemiology*, *54*(12), 1228–1237.

Shukla, P. K., & Deb, K. (2007). On finding multiple Pareto-optimal solutions using classical and evolutionary generating methods. *European Journal of Operational Research*, *181*(3), 1630–1652.

Sox, H. C., Higgins, M. C., & Owens, D. K. (2013). *Medical decision making* (2nd ed.). Chichester, UK: Wiley.

Stinnett, A. A., & Paltiel, A. (1996). Mathematical programming for the efficient allocation of health care resources. *Journal of Health Economics*, *15*(5), 641–653.

Taylor, F. C., Huffman, M., & Ebrahim, S. (2013). Statin therapy for primary prevention of cardiovascular disease. *JAMA*, *310*(22), 2451.

Tikkanen, E., Havulinna, A. S., Palotie, A., Salomaa, V., & Ripatti, S. (2013). Genetic risk prediction and a 2-stage risk screening strategy for coronary heart disease. *Arteriosclerosis, Thrombosis, and Vascular Biology*, *33*(9), 2261–2266.

Ubel, P. A., DeKay, M. L., Baron, J., & Asch, D. A. (1996). Cost-effectiveness analysis in a setting of budget constraints – is it equitable?*New England Journal of Medicine*, *334*(18), 1174–1177.

vanWalraven, C., Austin, P. C., Jennings, A., & Forster, A. J. (2009). Correlation between serial tests made disease probability estimates erroneous. *Journal of Clinical Epidemiology*, *62*(12), 1301–1305.

Vartiainen, E., Laatikainen, T., Salomaa, V., Jousilahti, P., Peltonen, M., & Puska, P. (2007). Sydäninfarkti- ja aivohalvausriskin arviointi FINRISKI-tutkimuksessa (The FINRISK FUNCTION: Estimation of the risk of coronary events and stroke in the Finnish population). *Suomen Lääkärilehti*, *48*(5), 4507–4513.

Weatherly, H., Drummond, M., Claxton, K., Cookson, R., Ferguson, B., Godfrey, C., et al. (2009). Methods for assessing the cost-effectiveness of public health interventions: Key challenges and recommendations. *Health Policy*, *93*(2), 85–92.

WHO (2016). Cardiovascular diseases (CVDs).

Wilson, P. W. F., D'Agostino, R. B., Levy, D., Belanger, A. M., Silbershatz, H., & Kannel, W. B. (1998). Prediction of coronary heart disease using risk factor categories. *Circulation*, *97*(18), 1837–1847.

Yaari, M. E., & Bar-Hillel, M. (1984). On dividing justly. *Social Choice and Welfare*, *1*(1), 1–24.

Zaric, G. S., & Brandeau, M. L. (2001). Optimal investment in a portfolio of HIV prevention programs. *Medical Decision Making*, *21*(5), 391–408.

APPENDIX A

Proposition A.1: *Let there be a single test with two possible results* $\{+, -\}$*, and assume that there are two treatment options* $\{T, N\}$ *referring to treatment or no treatment. Let*

$$\ell_J = \frac{(h(N|0) - h(T|0))p(+|0)}{(h(N|0) - h(T|0))p(+|0) + (h(T|1) - h(N|1))p(+|1)},$$

$$u_J = \frac{(h(N|0) - h(T|0))p(-|0)}{(h(N|0) - h(T|0))p(-|0) + (h(T|1) - h(N|1))p(-|1)}.$$

Then, the optimal segment-specific strategies at the highest expected cost level $j = J$ *are the same for both the utilitarian and egalitarian objective. Specifically, the optimal strategies are to*

1. treat patients with $p_0^i \geq u_J$,
2. not treat patients with $p_0^i \leq \ell_J$, and
3. test patients with $\ell_J < p_0^i < u_J$, after which those patients who get a positive result are treated and those who get a negative result are not.

Proof: At the highest cost level, the optimal overall strategies under both objectives are obtained by maximizing the expected segment-specific health outcomes. The expected benefits from treating/not treating a patient based on prior disease probability are $h(T|1)p_0^i + h(T|0)(1 - p_0^i)$ *and* $h(N|1)p_0^i + h(N|0)(1 - p_0^i)$*, respectively. Testing only makes sense, if different test results yield different decision recommendations, that is, a positive result results in treatment and a negative result in no treatment. Let us denote the probabilities of positive and negative test results by* p^+ *and* p^-*, respectively, where*

$$p^+ = p(+|1)p_0^i + p(+|0)(1 - p_0^i)$$
$$p^- = p(-|1)p_0^i + p(-|0)(1 - p_0^i).$$

Then, the expected health benefit from testing becomes

$$\begin{aligned}
&p^+\left(h(T|1)\frac{p(+|1)p_0^i}{p^+} + h(T|0)\frac{p(+|0)(1 - p_0^i)}{p^+}\right) \\
&\quad + p^-\left(h(N|1)\frac{p(-|1)p_0^i}{p^-} + h(N|0)\frac{p(-|0)(1 - p_0^i)}{p^-}\right) \\
&= h(T|1)p(+|1)p_0^i + h(T|0)p(+|0)(1 - p_0^i) + h(N|1)p(-|1)p_0 \\
&\quad + h(N|0)p(-|0)(1 - p_0^i) \\
&= p_0^i(h(T|1)p(+|1) + h(N|1)p(-|1) - h(T|0)p(+|0) \\
&\quad - h(N|0)p(-|0)) + h(T|0)p(+|0) + h(N|0)p(-|0).
\end{aligned}$$

Tests are only to be carried out if the additional information from testing is expected to result in a higher expected health outcome. This leads to the following two constraints on prior disease probability p_0^i:

$$\begin{aligned}
& p_0^i(h(T|1)p(+|1) + h(N|1)p(-|1) - h(T|0)p(+|0) - h(N|0)p(-|0)) \\
& \quad + h(T|0)p(+|0) + h(N|0)p(-|0) \\
& \quad > h(T|1)p_0^i + h(T|0)(1 - p_0^i) \\
\Leftrightarrow\ & p_0^i(p(-|1)(h(T|1) - h(N|1)) + p(-|0)(h(N|0) - h(T|0))) \\
& \quad < p(-|0)(h(N|0) - h(T|0)) \\
\Leftrightarrow\ & p_0^i < \frac{p(-|0)(h(N|0) - h(T|0))}{p(-|0)(h(N|0) - h(T|0)) + p(-|1)(h(T|1) - h(N|1))} \\
=\ & u_J
\end{aligned}$$

and

$$\begin{aligned}
& p_0^i(h(T|1)p(+|1) + h(N|1)p(-|1) - h(T|0)p(+|0) - h(N|0)p(-|0)) \\
& \quad + h(T|0)p(+|0) + h(N|0)p(-|0) \\
& \quad > h(N|1)p_0^i + h(N|0)(1 - p_0^i) \\
\Leftrightarrow\ & p_0^i(p(+|1)(h(T|1) - h(N|1)) + p(+|0)(h(N|0) - h(T|0))) \\
& \quad > p(+|0)(h(N|0) - h(T|0)) \\
\Leftrightarrow\ & p_0^i > \frac{p(+|0)(h(N|0) - h(T|0))}{p(+|0)(h(N|0) - h(T|0)) + p(+|1)(h(T|1) - h(N|1))} \\
=\ & \ell_J
\end{aligned}$$

□

Proposition A.2: *Let there be a single test with two possible results $\{+, -\}$, and assume that there are two treatment options $\{T, N\}$ referring to treatment or no treatment. Let*

$$\ell_1 = \frac{(c(T|0) - c(N|0))p(+|0) + c^{test}}{(c(T|0) - c(N|0))p(+|0) + (c(N|1) - c(T|1))p(+|1)}$$
$$u_1 = \frac{(c(T|0) - c(N|0))p(-|0) - c^{test}}{(c(T|0) - c(N|0))p(-|0) + (c(N|1) - c(T|1))p(-|1)}.$$

Then, the optimal segment-specific strategies at the lowest expected cost level $j = 1$ are the same for both the utilitarian and egalitarian objective. Specifically, the optimal strategies are to

1. treat patients with $p_0^i \geq u_1$,
2. not treat patients with $p_0^i \leq \ell_1$, and
3. Test patients with $\ell_1 < p_0^i < u_1$, after which those patients who get a positive result are treated and those who get a negative result are not.

Proof: At the lowest cost level, the optimal overall strategies under both objectives are obtained by minimizing the expected segment-specific costs. The expected costs from treating/not treating a patient based on prior disease probability are $c(T|1)p_0^i + c(T|0)(1-p_0^i)$ *and* $c(N|1)p_0^i + c(N|0)(1-p_0^i)$*, respectively. Testing only makes sense, if different test results yield different decision recommendations, that is, a positive result results in treatment and a negative result in no treatment. Hence, the expected cost from testing becomes*

$$\begin{aligned}
&p^+\left(c(T|1)\frac{p(+|1)p_0^i}{p^+} + c(T|0)\frac{p(+|0)(1-p_0^i)}{p^+}\right)\\
&\quad +p^-\left(c(N|1)\frac{p(-|1)p_0^i}{p^-} + c(N|0)\frac{p(-|0)(1-p_0^i)}{p^-}\right) + c^{test}\\
&= c(T|1)p(+|1)p_0^i + c(T|0)p(+|0)(1-p_0^i) + c(N|1)p(-|1)p_0\\
&\quad +c(N|0)p(-|0)(1-p_0^i) + c^{test}\\
&= p_0^i(c(T|1)p(+|1) + c(N|1)p(-|1) - c(T|0)p(+|0) - c(N|0)p(-|0))\\
&\quad +c(T|0)p(+|0) + c(N|0)p(-|0) + c^{test},
\end{aligned}$$

where p^+ *and* p^- *are the probabilities of positive and negative test results as in the proof of Proposition1.*

Tests are only to be carried out if the additional information from testing is expected to result in a lower expected cost. This leads to the following two constraints on prior disease probability p_0^i*:*

$$\begin{aligned}
&p_0^i(c(T|1)p(+|1) + c(N|1)p(-|1) - c(T|0)p(+|0) - c(N|0)p(-|0))\\
&\quad +c(T|0)p(+|0) + c(N|0)p(-|0) + c^{test} < c(T|1)p_0^i + c(T|0)(1-p_0^i)\\
\Leftrightarrow\ &p_0^i(p(-|1)(c(N|1) - c(T|1)) + p(-|0)(c(T|0)\\
&\quad -c(N|0))) < p(-|0)(c(T|0) - c(N|0)) - c^{test}\\
\Leftrightarrow\ &p_0^i < \frac{p(-|0)(c(T|0) - c(N|0)) - c^{test}}{p(-|0)(c(T|0) - c(N|0)) + p(-|1)(c(N|1) - c(T|1))}\\
=\ &u_1
\end{aligned}$$

and

$$\begin{aligned}
&p_0^i(c(T|1)p(+|1) + c(N|1)p(-|1) - c(T|0)p(+|0) - c(N|0)p(-|0))\\
&\quad +c(T|0)p(+|0) + c(N|0)p(-|0) + c^{test} < c(N|1)p_0^i + c(N|0)(1-p_0^i)\\
\Leftrightarrow\ &p_0^i(p(+|1)(c(N|1) - c(T|1)) + p(+|0)(c(T|0)\\
&\quad -c(N|0))) > p(+|0)(c(T|0) - c(N|0)) + c^{test}\\
\Leftrightarrow\ &p_0^i > \frac{p(+|0)(c(T|0) - c(N|0)) + c^{test}}{p(+|1)(c(N|1) - c(T|1)) + p(+|0)(c(T|0) - c(N|0))}\\
=\ &\ell_J
\end{aligned}$$

In both cases $j \in \{1, J\}$, the expected segment-specific health outcomes and costs are

$$h_{i,j} = \begin{cases} h(T|1)p_0^i + h(T|0)(1-p_0^i), & amp;\ p_0^i \geq u_j \\ [h(T|1)p(+|1) + h(N|1)p(-|1)]p_0^i & \\ \quad +[h(T|0)p(+|0) + h(N|0)p(-|0)](1-p_0^i) & amp;\ \ell_j < p_0^i < u_j \\ h(N|1)p_0^i + h(N|0)(1-p_0^i) & amp;\ p_0^i \leq \ell_j. \end{cases}$$

$$c_{i,j} = \begin{cases} c(T|1)p_0^i + c(T|0)(1-p_0^i), & amp;\ p_0^i \geq u_j \\ [c(T|1)p(+|1) + c(N|1)p(-|1)]p_0^i & \\ \quad +[c(T|0)p(+|0) + c(N|0)p(-|0)](1-p_0^i) + c^{test} & amp;\ \ell_j < p_0^i < u_j \\ c(N|1)p_0^i + c(N|0)(1-p_0^i) & amp;\ p_0^i \leq \ell_j, \end{cases}$$

□

Figure A.1 illustrates these segment-specific health outcomes and costs corresponding to highest and lowest cost levels, when the costs and health outcomes of (statin) treatment are as in the CHD example (see Table 1), and the probabilities of TRS test results are $p(+|1) = 0.85$ and $p(+|0) = 0.1$. The vertical lines show the lower and upper bounds for optimal testing regions at highest and lowest cost levels.

Given the distribution d_i of patients in each segment, the total expected health outcomes $H_j = \sum_i d_i h_{i,j}$ and costs $C_j = \sum_i d_i c_{i,j}$ corresponding to the lowest and highest cost levels $j \in \{1, J\}$ can be obtained in closed form.

APPENDIX B

Example of a Strategy Below the Convex Hull of the Pareto Frontier

Assume that there are two treatment options for a given disease: to treat or not to treat. Information about a patient's state of health can be obtained by carrying out one of three tests with two possible results (+,-) each. The first test is cheap but gives fairly inaccurate information, the second is significantly more expensive but only moderately more accurate, and the last is the most expensive but also the most accurate. The costs and health outcomes of the two treatment options, and the costs and accuracies of the tests are shown in Tables B.1 and B.2.

Figure B.1 shows the three Pareto optimal strategies (denoted by black markers) corresponding to the patient segment with a 70% initial probability of disease. The leftmost strategy corresponds to the use of test 1, the middle one to test 2, and the rightmost one to test 3. The strategy corresponding to test 2 lies below the convex hull of the Pareto frontier, denoted by the dashed line. This strategy would be dominated by a strategy in which 37.5% of patients in the segment selected at random would be tested with the cheap test 1, and the remaining 62.5% with the most expensive test 3 (denoted by a white marker). This strategy would have the same expected cost as the strategy corresponding to test 2, but would be expected to result in 0.197 additional QALYs.

Figure A.1: Optimal testing regions at lowest and highest cost levels, and the corresponding expected segment-specific health outcomes and costs.

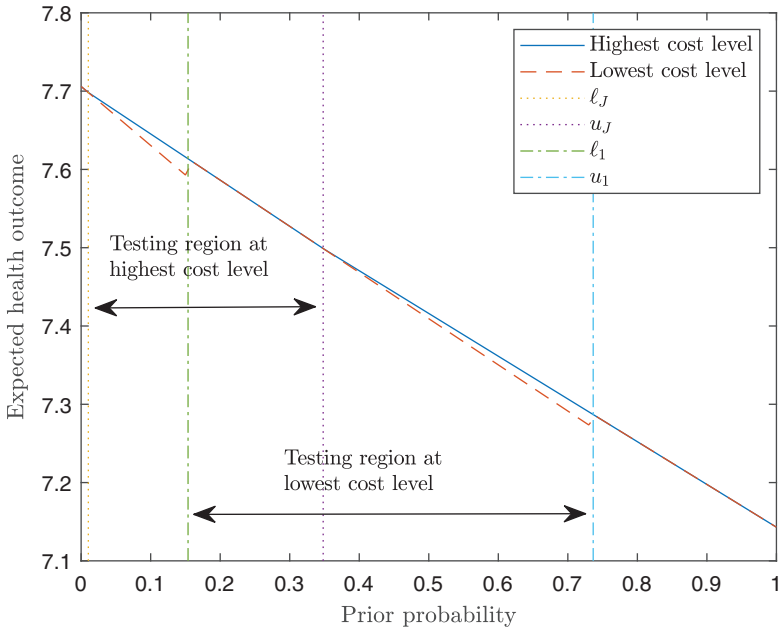

(a) Health outcomes

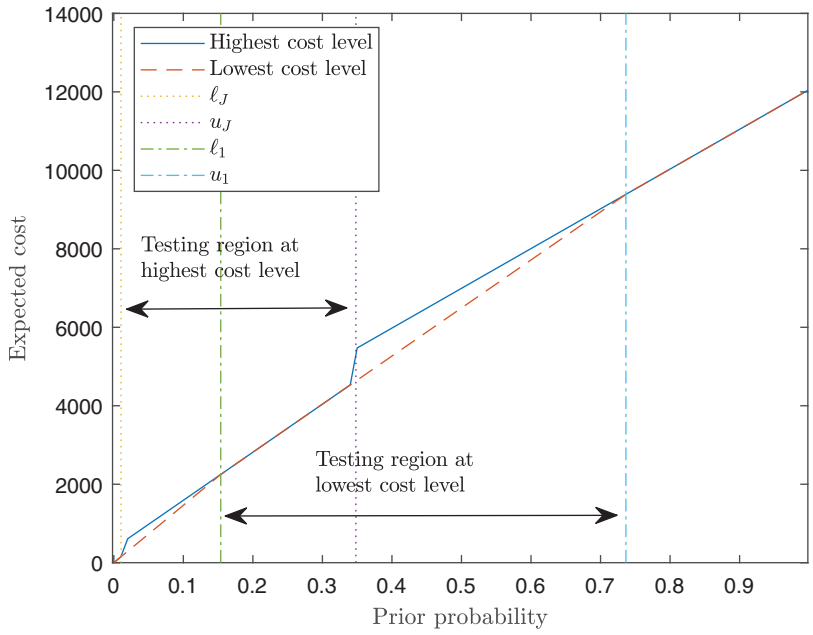

(a) Health outcomes

Table B.1: Costs and health outcomes.

	Treatment	No Treatment
Costs (€)		
Disease	8,000	10,000
No disease	2,000	0
Health outcomes (QALY)		
Disease	5	2
No disease	7	10

Table B.2: Probabilities of test results.

	+	−
Test 1 (50 €)		
Disease	0.7	0.3
No disease	0.2	0.8
Test 2 (200 €)		
Disease	0.75	0.25
No disease	0.15	0.85
Test 3 (500 €)		
Disease	0.9	0.1
No disease	0.05	0.95

Figure B.1: Pareto optimal strategies for the patient segment with 70% initial probability of disease (black markers). The convex hull of the Pareto frontier is denoted by dashed line. The randomized strategy on the convex hull with the same expected cost as the strategy corresponding to test 2 is denoted by a white marker.

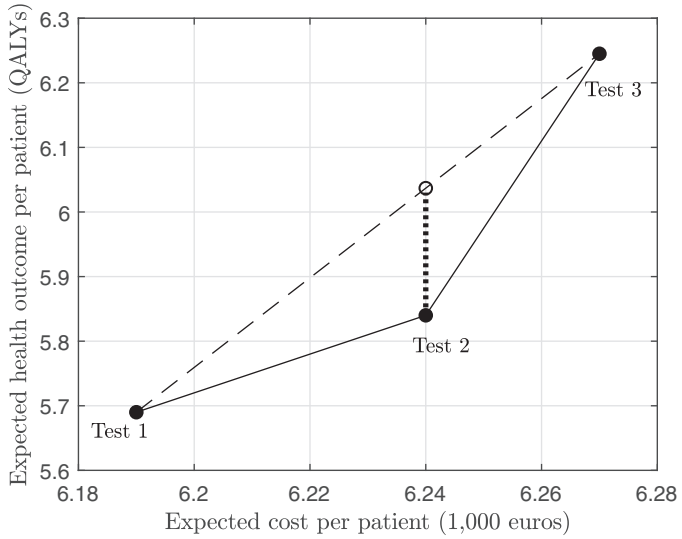

Dr. Yrjänä Hynninen received his PhD degree in Systems and Operations Research from the Aalto University in 2018. His research interests include the development of decision analytic methods to support cost-effective healthcare delivery. Currently, he is working as a Senior Manager and Head of Advanced Analytics at Finland-based Nordic Healthcare Group, a management consultancy specialized in planning and developing health and social services.

Dr. Eeva Vilkkumaa is an Assistant Professor of Management Science at the Department of Information and Service Management at Aalto University School of Business. Her research interests include the development of novel methods to support decision-making and resource allocation under uncertainty. Recent application areas include the cost-effective allocation of healthcare resources and scenario-based strategy development. Dr. Vilkkumaa is also the current president of the Finnish Operations Research Society.

Dr. Ahti Salo (http://sal.aalto.fi/ahti) is a professor of Systems Analysis at the Department of Mathematics and Systems Analysis at the Aalto University School of Science. He has worked extensively on the development of decision analytic methods for resource allocation, project selection, and risk management. He has received several awards for his research from the Decision Analysis Society of the Institute for Operations Research and the Management Sciences. In 2019, he won the MCDM Edgeworth-Pareto Award, the highest distinction of the International Society for Multiple Criteria Decision Making.