



This is an electronic reprint of the original article. This reprint may differ from the original in pagination and typographic detail.

Neuvonen, Lauri; Wildemeersch, Matthias; Vilkkumaa, Eeva

Supporting strategy selection in multiobjective decision problems under uncertainty and hidden requirements

Published in: European Journal of Operational Research

DOI: 10.1016/j.ejor.2022.09.036

Published: 16/05/2023

Document Version Publisher's PDF, also known as Version of record

Published under the following license: CC BY

Please cite the original version:

Neuvonen, L., Wildemeersch, M., & Vilkkumaa, E. (2023). Supporting strategy selection in multiobjective decision problems under uncertainty and hidden requirements. *European Journal of Operational Research*, *307*(1), 279-293. https://doi.org/10.1016/j.ejor.2022.09.036

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Contents lists available at ScienceDirect

# European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor



## Supporting strategy selection in multiobjective decision problems under uncertainty and hidden requirements



UROPEAN JOURNAL

## Lauri Neuvonen<sup>a,b,\*</sup>, Matthias Wildemeersch<sup>b</sup>, Eeva Vilkkumaa<sup>a</sup>

<sup>a</sup> Department of Information and Service Management, School of Business, Aalto University, Ekonominaukio 1, Espoo, 02150, Finland <sup>b</sup> International Institute for Applied Systems Analysis (IIASA), Schlossplatz 1, A-2361 Laxenburg, Austria

#### ARTICLE INFO

Article history: Received 8 September 2021 Accepted 25 September 2022 Available online 7 October 2022

MSC: 0000 1111

Keywords: Decision support systems Multiobjective optimization Robustness Pruning Implementability

## ABSTRACT

Decision-makers are often faced with multi-faceted problems that require making trade-offs between multiple, conflicting objectives under various uncertainties. The task is even more difficult when considering dynamic, non-linear processes and when the decisions themselves are complex, for instance in the case of selecting trajectories for multiple decision variables. These types of problems are often solved using multiobjective optimization (MOO). A typical problem in MOO is that the number of Pareto optimal solutions can be very large, whereby the selection process of a single preferred solution is cumbersome. Moreover, preference between model-based solutions may not be determined only by their objective function values, but also in terms of how robust and implementable these solutions are. In this paper, we develop a methodological framework to support the identification of a small but diverse set of robust Pareto optimal solutions. In particular, we eliminate non-robust solutions from the Pareto front and cluster the remaining solutions based on their similarity in the decision variable space. This enables a manageable visual inspection of the remaining solutions to compare them in terms of practical implementability. We illustrate the framework and its benefits by means of an epidemic control problem that minimizes deaths and economic impacts, and a screening program for colorectal cancer that minimizes cancer prevalence and costs. These examples highlight the general applicability of the framework for disparate types of decision problems and process models.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/)

## 1. Introduction

Decision-makers (DMs) are often faced with complex decisionmaking problems that require making trade-offs between multiple, conflicting objectives under various uncertainties. The task is even more difficult when considering dynamic, non-linear processes and when the decisions themselves are complex in form, for instance in the case of selecting profiles for multiple decision variables. Examples of such complex decision-making situations can be found in, e.g., epidemic control (Caulkins et al., 2020; 2021), pollution control (Lempert, Groves, Popper, & Bankes, 2006), water management (Kasprzyk, Nataraj, Reed, & Lempert, 2013), and production planning (Lin, Liu, Hao, & Jiang, 2016). These types of complex decision-making problems can be approached by (i) building a model to capture the complex dynamics between underlying processes and decision variables (see, e.g., Araz, Lant, Fowler, & Jehn, 2013; Klein, Dittus, Roberts, & Wilson, 1993; Miller et al.,

\* Corresponding author. *E-mail address:* lauri.neuvonen@aalto.fi (L. Neuvonen). 2005; Van Der Zee & Van Der Vorst, 2005), and (ii) optimizing the values of these decision variables by using multiobjective optimization (MOO) techniques (da Cruz, Cardoso, & Takahashi, 2011; Da Cruz, Cardoso, & Takahashi, 2009; Falke, Krengel, Meinerzhagen, & Schnettler, 2016; Meng, Lou, Peng, & Prybutok, 2017; Rangaiah, 2016; Deb, Pratap, Agarwal, & Meyarivan, 2002; Holzmann & Smith, 2018; Miettinen & Mäkelä, 1995). The result of an MOO process is a set of Pareto optimal solutions, which cannot be improved with respect to any objective without impairing performance on some other objective.

A common problem in MOO is that the number of Pareto optimal solutions is in many cases very large (Friedrich, Kroeger, & Neumann, 2011; Sudeng & Wattanapongsakorn, 2015; Wismans, Brands, Van Berkum, & Bliemer, 2014). Consequently, the task of directly selecting a single preferred solution from this set can be difficult. In particular, the objective function values and compliance with model constraints may not provide sufficient information for choosing the final solution suggested by the model but, rather, requirements related to the real-world implementability of this chosen solution must also be accounted for. This information

https://doi.org/10.1016/j.ejor.2022.09.036

0377-2217/© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/)



can, however, be hidden and therefore hard to elicit and quantify in such a manner that it could be included as a constraint in the optimization problem. Specifically, the DM may be able to judge the implementability of a given solution only after having visualized this solution and its relevant characteristics in comparison to those of other proposed solutions. But how can one offer the DM a small representative set of high-quality solutions, out of possibly hundreds of candidates so that such visual inspection can be carried out in a meaningful way?

This problem can be addressed by pruning the Pareto front, i.e., eliminating solutions from it based on one or more additional criteria. Pruning is typically done in the objective space. In their recent paper, Petchrompo, Wannakrairot, & Parlikad (2021) classify pruning methods in the objective space into three classes: preference-based methods, efficiency-based methods, and diversity-based methods. In preference-based methods, solutions are eliminated based on the DMs preferences with respect to the objective-function values or trade-offs between them. Salo & Hämäläinen (2010, 1995), for instance, suggest the use of multicriteria models with incompletely specified criterion weights to identify a subset of Pareto optimal solutions that are not dominated by any other feasible solution with respect to the DM's preference statements. Efficiency-based methods focus on identifying solutions that are efficient according to some pre-set indicators. Typically, methods in this class focus on identifying Pareto front knees, i.e., regions in which a small improvement in one objective would significantly worsen performance on at least one other objective (Sudeng & Wattanapongsakorn, 2016; Wismans et al., 2014). Diversity-based methods to prune the Pareto front in the objective space aim at identifying a diverse subset of solutions that together cover the range of different objective function values.

Diversity-based methods include clustering algorithms which, in effect, eliminate solutions whose objective function values are too close to one another. Taboada & Coit (2007), for instance, use an iterative k-means algorithm to find a small number of mutually dissimilar Pareto optimal solutions. Zio & Bazzo (2011) use a subtractive clustering approach for similar purposes. Yu, Zheng, Gao, & Yang (2017) integrate subtractive clustering with multi-criteria tournament decision and gain analysis methods to both maintain the shape of the Pareto front and consider the DM's preferences in choosing the final set of representative solutions. Petchrompo et al. (2021), apply k-medoids clustering to a portfolio asset management problem with the objective of reducing the full set of Pareto optimal solutions to a smaller representative set. Li, Liao, & Coit (2009) propose an approach in which self-organizing maps are first applied to cluster similar solutions together, after which data envelopment analysis is used to identify relatively efficient representative solutions within each cluster. The Pareto front could also be pruned in the decision space by, for example, clustering solutions based on the similarity of the corresponding decision variable profiles, and choosing a representative solution from each cluster. Pruning in the decision space would be particularly relevant from the perspective of implementability; yet, to our knowledge, these kinds of approaches have not been presented in the literature.

In addition to objective function values and diversity, many studies see robustness and risk considerations as important factors in pruning the Pareto front (Groetzner & Werner, 2021; Schöbel & Zhou-Kangas, 2021). By robustness we refer to the property of strategies to perform well in terms of objective function values and/or risk measures under different realizations of uncertain model parameters. Techniques for accommodating robustness considerations into multiobjective optimization problems have been presented by, e.g., Dellnitz & Witting (2009). A relatively recent contribution to accommodating robustness considerations into complex multiobjective decision-making problems is the many objective robust decision-making (MORDM) framework (Kasprzyk et al., 2013). MORDM seeks to combine the computational power of multiobjective evolutionary algorithms with robust decision-making techniques to help identify strategies that perform well across many different trajectories of the deep uncertainties affecting the underlying process. Interactive visual analytics are suggested to enable the exploration of trade-offs, robustness measures, and critical exogenous factors simultaneously. Such an interactive approach is beneficial in generating a deeper understanding for the DM about both the model and the decision-making problem at hand, but can be prohibitively time-consuming in situations requiring fast decisions. In addition to parametric uncertainty, the notion of robustness can be applied to other types of uncertainty such as variable uncertainty, which can be addressed by means of regularization robustness (Eichfelder, Krger, Schbel, & Eichfelder, 2015). These kinds of techniques can be beneficial in case there are uncertainties related to how accurately a chosen strategy can, in fact, be implemented. However, here we focus on parametric uncertainty.

In this paper, we propose a methodological framework for determining a small representative set of non-dominated, robust strategies to support DMs in finding implementable solutions to multiobjective decision-making problems under uncertainty. The framework is designed to be flexible with regard to modeling and optimization techniques, whereby it can be used in a wide range of problem settings. Importantly, the framework enables the utilization of existing models, which is beneficial in cases where it is possible to leverage the DM's trust towards certain models with which they are already familiar. In our proposed framework, we combine a process model with an MOO formulation to find the set of Pareto optimal solutions. Then, we prune the Pareto front (i) by eliminating solutions whose objective function values are nonrobust to small changes in the parameters of the process model, and (ii) by clustering solutions in the decision space to identify a small set of robust solutions that are sufficiently dissimilar from one another. Finally, we present effective visualizations of the remaining solutions to enable the DM to compare these solutions not only in terms of relevant objective function values and other performance metrics, but also their implementability. The generality of the framework is illustrated by means of two disparate case studies. The first case study analyzes how to control the spread of the coronavirus, minimizing impacts on both health and the economy for different strategy classes corresponding to lockdowns, mass testing, and combinations thereof. The second case study examines the optimal screening strategy for colorectal cancer, minimizing cancer prevalence and screening costs.

The contributions of this paper to the literature are threefold. First, the paper presents a widely applicable methodological framework for supporting decision-making in multiobjective problems in situations where the recommended solutions should be both robust as well as implementable. In particular, while implementability is a key consideration in various kinds of complex decision problems, no methods to support the identification of implementable solutions have previously been proposed. Second, in contrast to most existing methods, we prune the Pareto front in the decision space rather than in the objective space. This is particularly important in situations where the decisions correspond to selecting decision profiles, i.e. time-varying values for multiple decision variables, in which case the DM could be more interested in comparing a diverse set of such profiles rather than a set of similar profiles leading to different objective function values. Finally, our paper offers some insights into the effectiveness of different strategies for controlling the spread of the coronavirus and the prevalence of colorectal cancer.

The rest of the paper is structured as follows: In Section 2, we present the methodological framework. This framework is illustrated in detail in Sections 3 and 4 through two example case stud-

ies based on the COVID-19 epidemic and colorectal cancer screening. Finally, in Section 5, we present our conclusions and discuss the benefits and limitations of our proposed approach.

## 2. Methodological framework

#### 2.1. Framework motivation

We propose the framework to be a practical tool for decisionmaking especially under the following conditions: (i) When DMs are dealing with multiple objectives, which in the absence of preference information can result in a large set of alternative solutions. In this, clustering is a useful tool that allows DMs to evaluate a reduced set of diverse and representative non-dominated solutions. (ii) When there is considerable parametric uncertainty, which can be due to the absence of data to estimate the model parameters or to exogenous sources of uncertainty, for instance related to human behavior. The presence of model uncertainty requires us to study the robustness of non-dominated solutions. (iii) When there are hidden requirements related to the practical implementability of the solutions that the DM finds difficult to recognize or articulate as model constraints. This requires the visual inspection of complex decision vectors for which intuition on optimality and feasibility is elusive. (iv) When there exists time pressure to provide solutions expeditiously, preventing multiple interactions between DMs and system modelers. These time constraints can cause additional model uncertainty, stressing the need to analyze the robustness of non-dominated solutions.

## 2.2. Framework anatomy

The proposed methodological framework is summarized in Figure 1. In the first stage, a problem-specific model is designed and implemented to capture the most relevant problem characteristics and dynamics. This can be a model of any suitable type, e.g., an influence diagram, a simulation model, an agent-based model, or a dynamical systems model. In the second stage, MOO is used together with the problem-specific model to identify the set of Pareto optimal solutions. In some cases, it may be relevant to compute these sets for different classes of strategies, which results in multiple Pareto fronts. These strategy classes represent operating modes allowing for different decision options which can be in-dependently parameterised. An example of strategy classes in the context of controlling the COVID-19 epidemic is given in Section 3.

In the third stage of the methodology, strategies that are non-robust against parametric uncertainty are eliminated from the Pareto fronts. When using complex models, it is possible that integrating robustness considerations into the formulation of the optimization problem leads to numerical intractability. We therefore propose that non-robust solutions are eliminated postoptimization. In practice this can be done, for instance, by sampling sets of simulator parameters and running the simulator with these different parameter values for each Pareto optimal solution. This results in solution-specific distributions for the values of selected indicators, which permit the calculation of risk metrics (such as the probability of a given event or Conditional Value-at-Risk) that can be used to screen out non-robust solutions.

In the fourth stage of the process, multiple Pareto fronts are combined into one to eliminate solutions that are dominated by those belonging to another strategy class. This is done only after the robustness-based pruning, to avoid the risk of excluding robust solutions that are dominated by non-robust solutions. In the fifth stage, the solutions are clustered based on their similarity in the decision space to obtain a small but diverse set of representative solutions. Clustering is done only after robustness-based pruning

and combination of the Pareto fronts to make sure that no representative solutions need to be eliminated in subsequent stages due to inefficiency or non-robustness. By following this order, the number of clusters corresponds to the number of solutions that are to be subjected to closer visual inspection. In the sixth stage, effective visualizations are presented for the small set of remaining solutions to provide the DM a comprehensive view of each solution based on their objective function values, robustness, and implementability. The visualizations should present as many of the properties of the solutions relevant to the decision-making as possible, in a format that allows these solutions to be easily compared. In the context of the case studies in this paper, these include objective values, decision profiles and risk metrics. A careful inspection of such visualizations may reveal preferences or hidden requirements that the DM was unable to articulate at the beginning of the modeling process. Such an instance might require a reiteration of the framework starting at stage 1 or 2. Finally, the DM chooses one of the remaining solutions according to their preferences.

## 2.3. Framework applicability

The framework is applicable to many different decision-making problems, and as a consequence the framework is not specific about the optimization method, the clustering algorithm to prune the Pareto front, nor about the approach for robustness verification. The first two stages consist of setting up an MOO problem and computing the Pareto front. The framework is indifferent to the optimization approach, as long as the model later allows computing objective and risk measure values for different parameter samples for the non-dominated strategies. Stages 5 and 6 consist of clustering and visual analysis. Clustering is based only on decision variable values, and a suitable clustering algorithm needs to be identified to deal with the specific structure of the multidimensional decision vectors. Visual analysis of the solutions is a stage where some creativity is often needed. In general, the visualizations are highly dependent on the specific problem setting. That being said, there exist several effective techniques for visualizing, e.g., multidimensional objective and decision spaces, including interactive visualizations as well as projections to two- or three-dimensional spaces.

We have applied the framework, with different algorithmic choices within the stages, to analyze two multiobjective decision problems. The differences between the approaches used to tackle these two problems are profound, including process modeling methods, optimization algorithm structure and methods, and decision variable types. Yet, our proposed decision-support framework can be readily applied to both cases by selecting suitable optimization methods, clustering algorithms, and robustness metrics.

## 3. Case study 1: Epidemic control strategies for COVID-19

#### Problem context

We consider a hypothetical country with a population of 100 million people, and assume that there is a DM or a group of DMs in this country who wish to find an optimal strategy for controlling the spread of COVID-19. For this purpose, an epidemiological model has been built to capture the dynamics of the epidemic subject to different control strategies. Assume that the DM uses the epidemiological model to consider between six control strategy classes (representing a catalog of strategies followed by different countries) shown in Table 1: (i) mass testing with a capacity to carry out 3 million perfectly accurate tests per day, (ii) mass testing with a capacity to carry out 50 million tests per day with 85% sensitivity (i.e. probability  $q^+$  of correct identification of infected persons) and specificity (i.e. probability  $q^-$  of correct identification

European Journal of Operational Research 307 (2023) 279-293



Fig. 1. The decision support process for arriving at a small diverse set of non-dominated, robust strategies starting from a complex model and multiobjective optimization and followed by visual inspection and a decision.

 Table 1

 Strategy classes available to the DM. In all classes control strengths are selected at 30-day intervals.

Decision variables	$q^+$ & $q^-$	Max tests	Contact tracing
$ au_t$	100%	3M	no
$ au_t$	85%	50M	no
$\lambda_t^{LD}$	-	-	no
$\lambda_t^{LD}$	-	-	yes
$\lambda_t^{\text{LD}}, \tau_t$	100%	3M	no
$\lambda_t^{\mathrm{LD}}, \tau_t$	85%	50M	no
	$\begin{array}{c} \mbox{Decision variables} \\ \hline \tau_t \\ \tau_t \\ \lambda_t^{LD} \\ \lambda_t^{D} \\ \lambda_t^{D$	$\begin{array}{c c} \hline \text{Decision variables} & q^+ \& q^- \\ \hline \tau_t & 100\% \\ \hline \tau_t & 85\% \\ \lambda_t^{\text{LD}} & - \\ \lambda_t^{\text{D}} & - \\ \lambda_t^{\text{D}} & - \\ \lambda_t^{\text{D}}, \tau_t & 100\% \\ \lambda_t^{\text{D}}, \tau_t & 85\% \\ \end{array}$	$\begin{array}{c c c c c c c c c c c c c c c c c c c $

of non-infected persons), (iii) a lockdown strategy, (iv) a lockdown strategy with contact tracing, (v) a strategy that combines lockdowns with the possibility to carry out 3 million perfectly accurate tests per day, and (vi) a strategy that combines lockdowns with the possibility to carry out 50 million imperfect tests per day<sup>1</sup>. We model these different operating modes as separate strategy classes because we wish to evaluate their relative performance.

Within each strategy class, the DM is interested in finding Pareto optimal strategies that minimize cumulative COVID-19 related deaths and cumulative impact on the economy. A reasonable approximation for economic output is that it scales with the total

<sup>&</sup>lt;sup>1</sup> The numbers selected for this example case study are illustrative. Yet, they are realistic in the sense that high-sensitivity tests are limited by lab capacity whereas low-sensitivity antigen tests could be performed at much higher rates, as long as

the production of lateral flow tests can follow the demand. Hence, our results help generate qualitative, if not quantitative insights.

amount of active workers, i.e. those not in guarantine nor under lockdown (Berger, Herkenhoff, & Mongey, 2020). The decision variables for optimizing these strategies are the testing rate  $\tau_t$  (i.e., the share of the population that is tested at time t) and the meeting rate  $\lambda_t^{\text{LD}}$  for individuals affected by lockdown measures. In the epidemic model, we split the day into 14 time steps (cf. Berger et al., 2020). We allow the DM to adjust the decision profile at 30-day intervals<sup>2</sup> (29 for the first interval) starting from day 1 and continuing until day 570. In addition to minimizing COVID-19 related deaths and economic impact, the DM wants to ensure that the strategy that is ultimately selected will not lead to a risk of overloading the intensive care unit (ICU) capacity. Exceeding the ICU capacity significantly increases mortality and generates high social costs. Therefore, the DM wants to limit the probability of this happening given the uncertainties in key epidemic model parameter values. The DM is also interested in how much the suggested control strategies vary over time, as too much variance would lead to the strategies being unimplementable, both from a practical and a political viewpoint. However, the DM has not been able to formulate their preferences regarding implementability into explicit optimization constraints. In what follows, we show how each stage of our methodological framework could be carried out in the context of the case study.

#### Stage 1: Epidemic model

Building on the work of Berger et al. (2020), we develop a 12state compartmental model to capture the progression of the epidemic<sup>3</sup>. A visual representation of our model is shown in Figure 2, with a more detailed description given in S1 in the Supplementary Materials. The model describes in discrete time t the fractions of the population that belong to each of the 12 different compartments. This model extends the commonly used SEIR model (Berger et al., 2020; DeNegre, Myers, & Fefferman, 2020; IHME COVID-19 forecasting team, 2021), the compartments of which correspond to susceptible (S), exposed (E), infected (I) and recovered (R) members of the population. In our model, the 12 compartments correspond to non-infected (N), infected (I), or recovered (R) people who are either asymptomatic (A) or symptomatic (S). These people are either in guarantine (Q) or not in guarantine (NQ). The model accommodates incomplete information about whether an asymptomatic person is infected or not by differentiating between known and unknown states. In the visual representation, compartments corresponding to known states for asymptomatics are denoted by asterisks. By means of testing, individuals can be moved from the unknown to the known compartments. Compared to the existing work of Berger et al. (2020), our model enables the accommodation of the impacts of imperfect testing and contact tracing. Regarding imperfect testing, the model includes compartments corresponding to false positive (FP) and false negative (FN) test results. The details of modeling contact tracing are presented in S1.2, and variables related to contact tracing have the superscript index 'CT' in Figure 2.

The state equations capturing the dynamics of the compartments are compactly denoted by

$$X_{t+1} = \operatorname{SEIR}^+(X_t, X_{t-1}, \lambda_t^{\operatorname{LD}}, \tau_t)$$
(1)

where  $X_t$  represents the state of the system at time t, and SEIR<sup>+</sup> is a vector-valued function. A detailed description of these equations and the associated model parameters can be found in S1 of the Supplementary Material. The arguments of the state equations include the state at time t - 1 due to the accommodation of contract tracing (see S1.2).

## Stage 2: Multiobjective optimization

т

We consider two objectives in the optimization problem, i.e., the minimization of the number of COVID-19 related deaths  $D^{tot}$ and the minimization of economic costs  $\Omega$  over the considered time horizon. The economic cost at time *t* is represented by a relative loss of workforce and can be written as

$$C_{\rm e}(t,\boldsymbol{\lambda}^{\rm LD},\boldsymbol{\tau}) = 1 - \frac{\lambda_t^{\rm LD}}{\lambda} \frac{M_t^{\rm NQ}}{N} - \frac{\lambda_s^{\rm Q}}{\lambda} \frac{M_t^{\rm AQ}}{N}$$
(2)

where  $t \in \{0, ..., T\}$ ,  $\boldsymbol{\lambda}^{\text{LD}} = (\lambda_0^{\text{LD}}, ..., \lambda_T^{\text{LD}})$ ,  $\lambda_t^{\text{LD}} \in [0.5, 1.0]$ ,  $\boldsymbol{\tau} = (\tau_0, ..., \tau_T)$ ,  $\tau_t \in [0.0, 0.1]$ , *N* represents the total population, and  $M_t^{\text{NQ}}$  and  $M_t^{\text{AQ}}$  denote the number of individuals that are not quarantined and asymptomatic quarantined, respectively, at time *t*. Parameter  $\lambda$  is the default contact rate when no restrictions are in place.<sup>4</sup> Let  $\mathbf{u} = \{\boldsymbol{\lambda}^{\text{LD}}, \boldsymbol{\tau}\}$  belong to the set  $\mathcal{A}$  of admissible decision variables. Then, the optimization problem can be formulated as follows.

$$\min_{\mathbf{u}\in\mathcal{A}} \qquad D^{\text{tot}} = \sum_{t=0}^{I} D(t, \boldsymbol{\lambda}^{\text{LD}}, \boldsymbol{\tau}) + S_{D^{\text{tot}}}(T, \boldsymbol{\lambda}^{\text{LD}}, \boldsymbol{\tau})$$
(3)

$$\min_{\boldsymbol{u}\in\mathcal{A}} \qquad \Omega = \sum_{t=0}^{T} C_{e}(t, \boldsymbol{\lambda}^{\text{LD}}, \boldsymbol{\tau}) + S_{\Omega}(T, \boldsymbol{\lambda}^{\text{LD}}, \boldsymbol{\tau})$$
(4)

subject to  $X_{t+1} = \text{SEIR}^+(X_t, X_{t-1}, \lambda_t^{\text{LD}}, \tau_t)$  (5)

$$N_t^{\text{test}} \le N_t^{\text{test,max}}, \qquad t \in \{0, \dots, T\} \ \mathbf{u} \in \mathcal{A}$$
(6)

As to the terminal costs  $S_{\mathcal{X}}(T, \boldsymbol{\lambda}^{LD}, \boldsymbol{\tau})$  corresponding to each objective, we assume linear recovery after *T* over a recovery time  $\Delta T_{\text{rec}}$  (cf. Caulkins et al., 2021). The costs incurred after *T* over the economic recovery time  $\Delta T_{\text{rec}}$  can therefore be written as follows

$$S_{D}^{\text{tot}}(T, \boldsymbol{\lambda}^{\text{LD}}, \boldsymbol{\tau}) = \frac{\Delta T_{\text{rec}}}{2} D(T, \boldsymbol{\lambda}^{\text{LD}}, \boldsymbol{\tau}) ,$$
  
$$S_{\Omega}(T, \boldsymbol{\lambda}^{\text{LD}}, \boldsymbol{\tau}) = \frac{\Delta T_{\text{rec}}}{2} \left( 1 - \frac{\lambda_{T}^{\text{LD}}}{\lambda} \frac{M_{T}^{\text{NQ}}}{N} - \frac{\lambda^{\text{Q}}}{\lambda} \frac{M_{T}^{\text{Q}}}{N} \right) .$$
(7)

Although the assumption of linear recovery for all objectives is a simplification, it allows us to incorporate a consistent logic for all terminal costs. Moreover, the terminal cost is an aggregated cost, and therefore this cost term can capture different dynamics after terminal time *T*, depending on the value of  $\Delta T_{\text{rec}}$ .

We use the Non-dominated Sorting Genetic Algorithm II (NSGA-II, Deb et al., 2002) to find optimal decision profiles subject to the two objectives of minimizing deaths and economic impact. The details of the application of this algorithm are presented in S1.3. The NSGA-II algorithm is fairly commonly used to find Pareto optimal solutions for multiobjective decision problems but, due to its heuristic nature, cannot guarantee optimality. However, exact optimization algorithms cannot in practice be combined with complex simulation models due to computational issues. The choice of

<sup>&</sup>lt;sup>2</sup> The 30-day interval is a simplification but reflects real limitations that a DM might face: it is not possible to change national policies at a high rate due to issues of uptake, communication, and popular resistance to frequent policy changes.

<sup>&</sup>lt;sup>3</sup> Agent-based models (ABMs) can also be used to model epidemic dynamics (Aleta et al., 2020; Basurto, Dawid, Harting, Hepp, & Kohlweyer, 2020; Hoertel et al., 2020), but due to the high computational complexity associated with ABMs, they are, in practice, often hard to combine with optimization models in situations with a large set of feasible solutions and a large agent population.

<sup>&</sup>lt;sup>4</sup> As we do not consider a structured population model, the shares of individuals that are not quarantined and asymptomatic quarantined are the same in the workforce and in the total population.



Fig. 2. Extended SEIR compartmental epidemic model. Compartments where the state of the members is known are noted with \*. False information compartments are highlighted with red borders.

using a compartmental simulation model in conjunction with the NSGA-II algorithm is our way of balancing between sufficient realism in modeling epidemic dynamics and the desire to use optimization to find high-quality solutions to the complex decisionmaking problem. This choice enabled us to carry out all computations required for developing decision recommendations in our methodological framework within 4 days on a computing cluster using array runs.

The two-objective optimization problem (3)-(6) was solved separately for each of the six strategy classes. This resulted in six Pareto fronts shown in Figure 3a. Each of these fronts includes 60 solutions, totaling 360 solution candidates. The solutions are well spread between the extremes as a result of a diversity preference in the NSGA-II algorithm. The Pareto fronts corresponding to strategies without perfect testing coincide at the high-output extremity, corresponding to roughly 229 000 deaths (0.229% of the total population) and 99.1% economic output (i.e., 0.9% loss) compared to what would be obtained if there was no epidemic nor any control measures. These objective function values reflect the outcome of applying practically no control strategy at all. Perfect testing does not present any tradeoff, since perfect testing both improves economic output and lowers COVID-19 related deaths. However, the optimal strategy with perfect testing (green dot in Figure 3a corresponding to 128 000 deaths and 0.995 percent of economic output) is determined by the maximum number of tests. The combined strategy with perfect testing dominates the other strategies with economic outputs higher than approximately 0.95.

For lower values of economic output as well as deaths, the strategies with imperfect mass testing become dominating. At the other extreme of the fronts, there are solutions with slightly fewer than 1000 deaths and minimum output (0.5).

#### Stage 3: Elimination of non-robust solutions

Next, uncertainty in the epidemic parameter estimates is considered from the viewpoint of such uncertainties creating a risk of exceeding the ICU capacity when a Pareto optimal solution is implemented. To mitigate this risk, we eliminate solutions from the Pareto front that are non-robust in terms of resulting in a prohibitively high probability of exceeding the ICU capacity given small deviations from the parameters of the epidemic model. In particular, we exclude the solutions where the following probabilistic constraints cannot be met

$$\mathbb{P}\left[p \cdot \mathrm{ISQ}(t, \boldsymbol{\lambda}^{\mathrm{LD}}, \boldsymbol{\tau}) - Q_{\mathrm{hos}} > 0\right] \leq \varepsilon_{\mathrm{E}}, \qquad \forall t.$$
(8)

In these constraints, *p* is the probability that a symptomatic person requires ICU treatment, ISQ is the fraction of the population that is infected, symptomatic, and in quarantine (and can therefore be assumed to require hospital care),  $Q_{\text{hos}}$  is the hospital ICU capacity, and  $\varepsilon_{\text{E}}$  is the highest acceptable probability for exceeding the ICU capacity at any time *t*. Ideally, these constraints would already be included in the optimization task itself, but this would lead to numerical intractability.



(a) Stage 1: The Pareto fronts corresponding to all strategy classes.



(c) Stage 3: Remaining Pareto front clustered based on decision profiles in 8 clusters after removing dominated solutions from combined front.





(b) Stage 2: Remaining solutions after ICU overload risk cutoff at max 10% .



3) are representative (medoids of their clusters)

Fig. 3. Progression of the pruning process as seen in the objective space.

To estimate the probabilities on the left-hand side of Equations (8), we apply post-optimization sensitivity analysis on parameters  $R_0$  (basic reproduction rate) and  $\delta$  (symptom generation rate) that were found to have the strongest combined impact on model outcomes in simulations performed with strategies selected from the Pareto fronts (see section S1.4 in the Supplementary Materials for details). Specifically, we generated 10 000 parameter value samples from the distributions of these two parameters (see Table S4 for details), and estimated the probability of exceeding the ICU capacity for each strategy as the share of samples for which  $p \cdot ISQ(t, \lambda^{LD}, \tau, \phi) \ge 0$  for any t, where  $\phi = (R_0, \delta)$  is an individual sample. The solutions for which this share was higher than  $\varepsilon_E = 10\%$  were eliminated from the Pareto front. The remaining 230 robust solutions are shown in Figure 3b.

#### Stage 4: Combining the Pareto fronts

After robustness-based pruning, the Pareto fronts corresponding to different strategy classes are combined to eliminate solutions that are dominated by other solutions from a different class. In our case study, the combination of the Pareto fronts resulted in a set of 74 remaining solutions, which are illustrated in Figure 3c.

#### Stage 5: Clustering the remaining Pareto optimal solutions

Next, we cluster the remaining solutions in the decision space to obtain a small number of representative solutions that together cover a diverse set of decision profiles. The k-medoids algorithm is suitable for when solutions correspond to correlated decision profiles over time. The k-medoids algorithm minimizes a distance metric with respect to a representative object in the cluster (Kaufmann & Rousseeuw, 1987). As we aim to cluster solution profiles with similar decision profiles, we use the Pearson distance to group highly correlated solution profiles in a single cluster. Strategies from different strategy classes are set to feature a large Pearson distance between them, resulting in no mixing of strategy classes within clusters. Let  $u^{*i}$  and  $u^{*j}$  be the control vectors of two Pareto optimal solutions and let  $r_{u^{*i}u^{*j}}$  be the Pearson correlation coefficient. Then, we can define the Pearson distance as

$$d_{u^{*i}u^{*j}} \doteq \sqrt{\frac{1}{2}(1 - r_{u^{*i}u^{*j}})}.$$
(9)

The clustering algorithm selects one of the members of a cluster as a medoid, i.e., a representative solution for the considered cluster.

We cluster the remaining 74 solutions into 10 clusters based on strategy class and Pearson distance. The results of the clustering



**Fig. 4.** Summary of results after all pruning steps for visual inspection. In the strategy graphs, red lines depict testing strategies and blue lines lockdown strategies. Solutions selected by the DM for final comparison are highlighted with red dots in the strategy graphs and red circles in the Pareto front. The maximum risk level (%) over the considered time window is indicated next to each Pareto point. (MT = mass testing, LD = lockdown, CT = contact tracing, s&s = sensitivity & specificity, combo = combination strategy with lockdown and testing).

are shown in Figure 3c, where solutions corresponding to different clusters are depicted by different colours. Even though clustering is done in the decision space, we can see that there is little overlap between clusters in the objective space. The 10 remaining medoids are shown in Figure 3d.

## Stage 6: Visual inspection of representative solutions

At this stage, the number of Pareto optimal solutions has been pruned from 360 to 10, corresponding to a 97% reduction from the original set. The remaining solutions represent robust, nondominated, and relatively dissimilar control strategies, and hence they are attractive candidates to be chosen as the final solution. Due to the small number of these solutions, visual inspection can be applied to compare them with one another. An example of an effective visualization is shown in Figure 4, where red lines in the small, solution-specific figures depict the testing rate and blue lines depict the strength of lockdown at each time period.

The DM can now make judgements on the implementability of the solutions based on 1) their decision profiles and 2) the progression of the epidemic under the remaining strategies (see Figure S2 in the Supplementary material section S1.5). Assume that the DM holds large temporal fluctuations in the decision profiles as unimplementable, and is interested in further investigating solutions at the extremes of the remaining Pareto front and in the knee region (where a small improvement in one objective leads to a significant deterioration in the other). Based on visual inspection of the medoids presented in Figure 4, such a DM might select the first, seventh, and tenth medoid (counting clockwise from the bottom left) for further consideration. The characteristics of these three solutions are shown in Table 2. Strategy 1, where a relatively strong lockdown is imposed throughout the entire time horizon, leads to a low number of deaths, but also to the economic output being reduced to one half. Here, the risk of ICU overload is small (only 1.0%). Strategy 3, where a combination of relatively mild lockdowns and low testing rates with perfectly accurate tests is used, represents the other extreme: economic output is reduced by only 3%, but the number of deaths is approximately 86 000. Strategy 2, which corresponds to mass testing with imperfect tests, can be seen as a compromise between the two extremes: the economic output is only 4 percentage points lower than with strategy 3, but the number of deaths is cut down by a factor of four approximately.

#### Strategy implications and insights for controlling COVID-19

During the decision support process and related analyses, several insights were gained into the differences between alternative strategies and their implications from the viewpoint of the epidemic's progression. These insights would be hard to obtain without the use of optimization and the ability to ultimately focus on a small set of representative solutions that together cover the main characteristics of different strategies. Yet, the nature of these insights should be considered qualitative instead of quantitative as some of the underlying model assumptions have not been validated by experts or based on peer-reviewed literature.

The first insight is related to the positioning of strategies of different class on the Pareto front (cf. Figure 3a). In particular, strate-

#### Table 2

Remaining non dominated solutions after risk and clustering based pruning and assumed DM preferences.

Strategy #	Strategy class	Deaths	Economic output	ICU overload risk
1	Lockdown	920	0.5	1.0%
2	Mass testing (sens & spec 85%)	22238	0.93	6.1%
3	Combination (sens & spec 100%)	86080	0.97	8.6%

gies at the low death rate, low economic output end of the Pareto front use lockdowns, whereas those resulting in high death rates and high economic output correspond to testing with perfectly accurate tests (at a restricted capacity). Most strategies of the latter class are nevertheless eliminated in stage 3 of the framework based on robustness considerations, since a high death rate is connected to a high risk of overloading the ICU capacity. In the knee area, Pareto optimal strategies use mass testing with imperfect tests. Specifically, all of these strategies use high volumes of testing in the beginning, which significantly slows down the epidemic. The longer these high testing volumes are maintained, the fewer COVID-19 related deaths there will be.

The second insight concerns a so-called 'lockdown by testing' effect related to strategies that use mass testing with imperfect tests. This effect stems from the large number of false positive results in these strategies, which lead to a large portion of the non-infected population being quarantined (cf. Figure S2 in the Supplementary material). Thus, direct lockdown measures may become obsolete. In fact, pure lockdown strategies are dominated by those that utilize imperfect tests as long as there is enough testing capacity. Nevertheless, issues related to accountability and public trust may render strategies utilizing the 'lockdown by testing' effect unimplementable in real-life situations.

Our third insight is that all strategies which remain in the Pareto front after robustness-based pruning incorporate lockdowns in some way, either directly or indirectly through the 'lockdown by testing' effect (cf. Figure 4). To mitigate the risk of ICU overloads (and, thereby, high death rates) with strategies utilizing perfect tests and no lockdowns, a large testing capacity would be required.

#### 4. Case study 2: Colorectal cancer screening

#### Problem context

In Finland, colorectal cancer (CRC) is a crucial concern for public health. Incidence rates have increased over the last decades, and as of 2017, CRC is the sixth most common cause of death. Screening is an effective method to catch and treat potential cancers in their early stages, thus improving the prognosis of CRC patients significantly. Screening for CRC is a multi-period process, where participants are screened in, e.g., 2-year intervals. The Finnish program utilizes faecal immunochemical testing (FIT) to first filter participants for further examination. Then, those whose FIT result exceeds a cut-off level are invited for a colonoscopy (Finnish colorectal cancer screening expert groups, 2021). During the colonoscopy, a visual inspection of the colon is performed, and detected abnormalities in the bowel are recorded, sampled and possibly removed. These abnormalities include small benign growths (polyps), different sizes and stages of larger growths (adenomas), and cancerous growths (cancers). The samples are then studied to diagnose potentially cancerous growth, and depending on the results, the correct treatment is performed.

In this example case study, we apply the proposed methodological framework to help improve the current Finnish CRC screening program. For the purposes of properly demonstrating each stage of our framework, we modify the actual case study presented in Neuvonen et al. (2022), where the model was built in collaboration with practitioners. In the current program, the FIT cut-off level is assumed to be the same for both sexes and all age groups, which may result in a suboptimal allocation of colonoscopy resources. Moreover, due to uncertainties related to FIT results and participation in the program, there is a risk of exceeding the existing capacity of 18 000 colonoscopies, which could lead to failures in carrying out the program or costly rearrangements. To accommodate these considerations, we develop a multistage optimization approach based on multiobjective influence diagrams. Specifically, we optimize the age- and sex-specific FIT cut-off levels for a fiveperiod screening program in view of minimizing both expected cancer prevalence and expected costs in the total target population with given colonoscopy resources. Moreover, to mitigate the risk of exceeding colonoscopy capacity, we aim to eliminate solutions that have a higher than 10% chance of exceeding this capacity. We further assume that the DM is interested in ensuring that the strategies make sense from a behavioral perspective, but has no quantitative metrics for formally including such considerations.

A schematic description of the multistage optimization approach is shown in Figure 5. All relevant details of this approach can be found in Neuvonen et al. (2022). In phase 1 (performed separately for both sexes), influence diagrams (IDs) are used to capture how abnormal bowel states are found and costs generated in a given period as a function of screening decisions. An ID can, under certain assumptions, be transformed into a decision tree and solved accordingly (Howard & Matheson, 2005). The case-studyspecific IDs will be presented in more detail in the description of stage 1 of the methodological framework. Within a single period, the Modified Augmented Weighted Tchebychev (MAWT) algorithm (Holzmann & Smith, 2018) is used to optimize the ID for a given target segment defined by sex (g) and age group (k), given starting prevalences  $\psi_{g,k,b}$  for different bowel states *b* in this segment. This produces a set of Pareto optimal solutions with respect to minimizing expected costs and maximizing the expected number of detected abnormal bowel states. Corresponding to each such solution, the starting prevalences of different bowel states for the next period are produced through prevalence update rules that also account for the effects of aging. This process is repeated until all target segments have been optimized for all possible decision histories, i.e., strategies until that period. In phase 2, the Pareto fronts of these sex-specific strategies are combined into full strategies, after which dominated strategies are removed. The number of strategies available is nearly 26 billion, which, together with chance events and parametric uncertainties, makes this a numerically complex decision-making problem.

#### Stage 1: Cancer screening model

The cancer screening process for an individual member of the population in a given period and with a given starting prevalence is modeled by an ID. This ID, shown in Figure 5, is based on the current Finnish CRC screening programme. The ID contains decision nodes (squares), chance nodes (circles), and utility nodes (diamonds). An instance where exactly one value is realized for each chance and decision node is called a *path* through the ID. A choice of a combination of values for the decision nodes (given the real-



SCREENING OPTIMIZATION ALGORITHM

Fig. 5. Schematic description of the multistage optimization approach to solve the CRC screening optimization problem.

izations of chance nodes preceding these decision nodes, in case there are any) is called a *strategy*. Thus, a strategy is a collection of paths corresponding to all possible realizations of chance nodes that can be attained given the choice of values for the decision nodes.

In Figure 5, the decision nodes 1, 2, and 3 correspond to decisions about 1) what FIT cut-off value to use to select participants for a colonoscopy in the given target segment, 2) whether to use an incentive to boost participation rate among invitees in this segment (specifically, we assume that an incentive worth 10 euros halves the number of non-returned samples) and 3) whether to invite the target segment to the screening program. Let us denote by  $s_i \in S_i$  the alternatives for decision *j*. Here, the sets  $S_i$  of such alternatives are  $S_2 = S_3 = \{yes, no\}$  and  $S_1 = \{10, 20, 30, 40, 50\} \mu g$ Hg/g of blood in the stool sample. Let (g, k) with  $g \in \mathcal{G} = \{F, M\}, k \in$  $\{1, 2, 3, 4, 5\}$  be the target segment, where F and M refer to female and male, respectively, and numbers k to the screening period. These periods correspond to a participant's age in 2-year intervals so that period 1 refers to 60-year-olds, period 2 to 62-yearolds etc. The decisions regarding which alternatives  $s_i$  to select for each target segment are modeled as binary decision variables  $z_{g,k}(s_i) \in \{0, 1\}$  so that  $z_{g,k}(s_i) = 1$  if and only if alternative  $s_i$  is selected for segment (g, k). A full screening strategy Z is a collection of these decision variables for all target segments and can be written as  $Z = \bigcup_{g,k} Z_{g,k}$  where the segment-specific strategies are defined as  $Z_{g,k} = \bigcup_{j \in D} z_{g,k}(s_j)$ ,  $\forall g, k$ .

Chance nodes correspond to returning a FIT sample, FIT results, continued participation, the discovery of polyps and polypectomy in a colonosopy, and adverse effects from the colonoscopy. In particular, nodes 5 and 8 reveal information about the participant's bowel state. Here, we assume that the bowel state  $b \in \mathcal{B} = \{N, B, L, R\}$  of a participant can be Normal (N), Benign growth (B), Large growth (L) or CRC (R). These bowel states are reflected at the population level by prevalences  $\psi_{g,k,b} = N_{g,k,b}/N_{g,k}$ , where  $N_{g,k,b}$  is the number of participants in segment (g, k) with bowel state b, and  $N_{g,k}$  is the total number of participants in segment (g, k). The starting prevalences  $\psi_{F,1,b}$ ,  $\psi_{M,1,b}$ ,  $\forall b \in \mathcal{B}$  as well as the segment-specific conditional probabilities for chance nodes are obtained from the literature.

The progression of the prevalences of different bowel states are affected by two factors: natural progression and screening. The natural progression of colorectal cancer is reflected by transition probabilities  $\mathbb{T}_{b,b'}^{g,k}$  between bowel states  $b, b' \in \mathcal{B}$ . We assume that this progression follows the adenoma-carcinoma sequence, meaning the transition through bowel states can be represented by a linear recurrence relation. Screening, on the other hand, helps de-

crease the prevalences of abnormal bowel states in the population, depending on the selected FIT cut-off level and incentive. In particular, we assume that any benign growth, large growth, or CRC found during the screening pathway is removed and that the bowel returns to a normal state. Taken together, the starting prevalences in period k + 1 can be computed using the following difference equations:

$$\psi_{g,k+1,B}(Z_{g,k}) = (\psi_{g,k,B} - \tilde{\psi}_{g,k,B}(Z_{g,k}))(1 - \mathbb{T}_{B,L}^{g,k}) + \psi_{g,k,N}\mathbb{T}_{N,B}^{g,k}$$
(10)

$$\psi_{g,k+1,L}(Z_{g,k}) = (\psi_{g,k,L} - \tilde{\psi}_{g,k,L}(Z_{g,k}))(1 - \mathbb{T}_{L,R}^{g,k}) + (\psi_{g,k,B} - \tilde{\psi}_{g,k,B}(Z_{g,k}))\mathbb{T}_{B,L}^{g,k}$$
(11)

$$\psi_{g,k+1,R}(Z_{g,k}) = \psi_{g,k,R} - \tilde{\psi}_{g,k,R}(Z_{g,k}) + (\psi_{g,k,L} - \tilde{\psi}_{g,k,L}(Z_{g,k})) \mathbb{T}_{L,R}^{g,k}$$
(12)

$$\psi_{g,k+1,N}(Z_{g,k}) = 1 - \sum_{b \in \{B,L,R\}} \psi_{g,k+1,b}(Z_{g,k}) \quad , \tag{13}$$

where  $\tilde{\psi}_{g,k,B}(Z_{g,k})$  stands for the fraction of participants found to have bowel state *b* in period *k* as a result of applying screening strategy  $Z_{g,k}$ .

Finally, the utility nodes describe the outcomes of a given path through the ID. Here, the 'Cost' node sums all costs accrued within the path, whereas nodes 'Found R', 'Found L' and 'Found B' obtain a value of 1 if CRC, large adenoma, or benign growth are found along the path, respectively, and a value of 0 otherwise. Node 'P Col.' collects information about whether a colonoscopy was carried out within the path. A strategy is associated with a distribution of outcomes for each utility node, where these distributions are determined by the probabilities of the paths corresponding to this strategy. These distributions can be used to formulate objectives or constraints for the optimization problem.

## Stage 2: Multiobjective optimization

We assume that the DM wants to minimize the combined expected prevalence of colorectal cancer across all target segments  $\Psi_{\rm R}(Z) = \sum_{g,k} N_{g,k,{\rm R}}(Z) / \sum_{g,k} N_{g,k}$  as well as the total expected costs  $\Omega(Z)$  so that the expected total number of colonoscopies  $N_{Col}(Z)$  does not exceed the capacity of 18 000. The multi-period optimization problem can thus be formulated as follows:

$$\min_{Z} \quad \Psi_{R}(Z) \tag{14}$$

$$\min_{Z} \quad \Omega(Z) \tag{15}$$

s.t. 
$$N_{Col}(Z) \le 18000.$$
 (16)

As illustrated in Figure 5, the optimization of strategies for the complete screening program is carried out in two phases. In phase 1, the MAWT algorithm (Holzmann & Smith, 2018) is used for finding the set of Pareto optimal solutions for each segment (g, k) in view of minimizing expected costs and maximizing the probability of finding abnormal bowel states (B, L, R). The MAWT algorithm was chosen as it helps generate the complete Pareto front for a discrete multiobjective optimization problem with any number of objectives relatively efficiently compared to other existing methods. Specifically, phase 1 of the screening optimization algorithm starts by identifying the Pareto optimal solutions for the first period for both sexes based on starting prevalences that have been

estimated from data. Subsequently, for each of these solutions, updated prevalences are computed using Equations (10) -(13) to generate a list of possible starting prevalence vectors for the next period. This process is repeated in the next period for each possible starting prevalence vector and continued until the last screening period. This results in a strategy tree, where each branch corresponds to a complete screening strategy. Dominated strategies as well as those for which the expected total number of colonoscopies exceeds the capacity of 18 000 are cut out during the process. In phase 2, a single Pareto front is produced for the multiperiod optimization problem defined in (14)-(16) by first creating all possible combinations of the female and male strategies, and the expected cancer prevalences and costs corresponding to these strategy combinations. Then, the Pareto front is obtained by removing all dominated strategy combinations from this set. The result is depicted in Figure 6a. At this stage, the front consists of 181 solutions.

## Stage 3: Elimination of non-robust solutions

In stage 3, the aim is to eliminate those solutions from the Pareto front for which the probability of exceeding the capacity for colonoscopies is, due to uncertainties in the model parameters, higher than 10%. The most important parameters in this respect are the probability of returning a FIT sample (cf. chance node 4 in the ID) and the sensitivity (i.e., true positive rate) and specificity (i.e., true negative rate) values of the FIT with a given cut-off level determining a positive result (cf. chance node 5 in the ID). This is because the probability of returning the sample was estimated from previous trials where a different sampling method was used, and because the FIT sensitivity and specificity estimates were based on literature regarding non-Finnish populations. Furthermore, the probability of a colonoscopy being performed depends on these parameters through equations defining the ID (see Section S2.1 in the Supplementary Material for details). We, nevertheless, assume that the parametric uncertainty related to continuing in the program after a positive FIT result is low because the probability of continuation can be reliably estimated based on previous screening trials in the Finnish population. Furthermore, we assume that the starting prevalences and transition probabilities have been accurately estimated, whereby uncertainty related to prevalences produced by the model in subsequent periods can be excluded from this robustness analysis as well

To estimate the probability of exceeding the capacity of 18 000 colonoscopies, we created 10 000 samples of the chosen parameters using the following distributions: The FIT return rate is assumed to follow a truncated normal distribution with expected value at the original parameter value estimate and a standard deviation of 10% of the expected value. FIT sensitivity is sampled as a proportional deviation from the original estimate, where the deviation is assumed to follow a truncated normal distribution between -1 and 1, with a standard deviation of 0.1 and expected value of 0. The deviation is the same for all FIT sensitivity parameters per sample, i.e., all FIT sensitivity parameters are assumed to be perfectly correlated. FIT specificity samples are modeled in a similar fashion. Their deviations are assumed to be independent from those of the sensitivity parameter. An estimate for the probability of exceeding the maximum colonoscopy capacity for a given strategy Z is obtained as the ratio of samples in which this capacity was exceeded to the total number of samples. The solutions for which this estimate was higher than 10% were removed, which led to the low-risk Pareto front in Figure 6b. The number of solutions remaining in the Pareto front after robustness-based pruning was 137.



(a) Stage 1: The Pareto front after phase 2 of the screening optimization algorithm.



(c) Stage 3: Remaining Pareto front clustered based on decision profiles in 10 clusters after removing dominated solutions from combined front.





(b) Stage 2: Remaining solutions after colonoscopy overdem and risk cutoff at max 10%.



 (d) Stage 4: centrold solutions of the clusters. These solutions are now 1) non-dominated,
 2) robust, and 3) representative (centroids of their clusters)

Stage 4: Combining the Pareto fronts

Since we did not consider different strategy classes in this case study, stage 4 can be omitted.

#### Stage 5: Clustering the remaining Pareto optimal solutions

Strategies in this case study can be considered categorical by type, in that two out of three decision variables (whether to invite the target segment to the program and whether to use an incentive to boost participation) are binary (yes / no), and the third variable (the FIT cut-off value) only has five possible values. Hence, we apply the k-modes approach (Chaturvedi, Green, & Caroll, 2001) to cluster the remaining 137 strategies in the decision space into 10 clusters. The k-modes approach uses a similarity measure that counts matches in categories between the members of a cluster. Clustering was performed using the 'kmodes' (v.0.11.1) package for Python (de Vos, 2015–2021). The effect of clustering in the objective space can be seen in Figure 6c.

## Stage 6: Visual inspection of the representative solutions.

The representative solutions for each of the 10 clusters are presented along with their objective values in Figure 7. The 10 remaining solutions represent 5.5% of the initial 181 solutions. The probability of exceeding the capacity for colonoscopies is indicated next to the marker for the representative solution in the objective space. It can be seen that solutions for which this probability is highest are mostly found in the high-cost end of the spectrum. This, however, is to be expected as a large portion of possible costs are related to the colonoscopy operation, and treatment and adverse effects depending on its outcome and results.

The DM can use Figure 7 to review the representative strategies and visually assess their implementability. In this particular case, the implementability of a strategy could be judged by considering whether the use of incentives within a strategy makes sense from a behavioral perspective. For example, the two incentivized strategies could be considered unimplementable due to incentives being used for some invited age groups but not all. This kind of an incentive structure could prove to be difficult to communicate to program invitees and, at worst, lead to the youngest age groups not participating at all. The rest of the strategies could be considered equally implementable and, coincidentally, equally robust in that each of them also has a 0.0% probability of exceeding the colonoscopy capacity. Among these strategies the DM might then choose, e.g., the one that minimizes the expected total can-

Fig. 6. Progression of the pruning process as seen in the objective space.



**Fig. 7.** Summary of CRC case study results after all pruning steps for visual inspection. In the strategy graphs, red triangles depict the selected strategy for females and blue triangles for males. The solution potentially selected by the DM is highlighted with a red dot in the strategy graph and a red circle in the Pareto front. The risk of exceeding the maximum colonoscopy limit (%) is indicated next to each Pareto point.

cer prevalence in the population (i.e., the strategy corresponding to the lower-left corner strategy graph).

## Insights from the cancer screening case study.

Several insights were gained as a result of the application of the framework to this problem. As in the epidemic case study, these insights should be considered qualitative by nature. The first insight is that in most cases, screening with a low FIT cut-off level is preferred (cf. strategy graphs in Figure 7). Specifically, the cut-off levels used in most representative strategies vary between 10-20  $\mu$ g Hg/g, whereas the cut-off level used in the current Finnish screening strategy is 25  $\mu$ g Hg/g. A possible explanation for the preference for low cut-off levels is that the use of a higher cut-off level decreases the number of performed colonoscopies, whereby a larger share of abnormal bowel states could go undetected. This, in turn, would lead to lower health benefits. Nevertheless, the specific reasons for why the suggested cut-off levels differ from current practice merits a more thorough investigation.

Our second insight is that for females, the lowest FIT cut-off level is used for nearly all age groups that are invited to the screening program. On the other hand, males in many representative strategies should start with a higher cut-off level which is then decreased in the older age groups. This can be interpreted to mean that for females the improved hit rate in screening resulting from the age-related increase in cancer prevalence does not cancel out the benefits gained from catching earlier stages in cancer development early on, whereas for males the situation might be reversed. A third insight is that, under these assumptions, incentives would not be very cost-effective, and are not used except for males in the high-cost end of the Pareto front.

#### 5. Conclusions and discussion

In this paper, we have developed a methodological framework to support the identification of a small but diverse set of robust Pareto optimal solutions to complex, non-linear decision-making problems. The main benefit of our framework is that it helps prune a possibly very large set of Pareto optimal solutions to a handful of robust, non-dominated solutions that represent a diverse set of decision alternatives. The small number of such solutions enables their thorough visual inspection, which can help make judgments about the relative characteristics and implementability of these solutions in view of practical and political criteria that are not easily converted into constraints for the optimization model. Visual inspection in the decision space can be particularly helpful when the number of objectives is large so that comparisons in the objective space become difficult. Another benefit of our framework is that between initial problem structuring and the visual inspection of the remaining solutions, no interaction with the DM is necessarily required. This can be useful in cases where time constraints are a major issue. If time permits, more interaction can of course be introduced to different stages of the framework to enhance learning and produce a sense of ownership of both the model and the decision recommendations (Franco & Montibeller, 2010). Such interaction could, for instance, include the iteration of the optimization model to include constraints related to implementability requirements that were hidden at the start of the decision process, but ultimately discovered through visual inspection. Finally, our framework is generic in that it can be applied to various types of process and optimization models (including black box models), as long as the computational burden required for solving the Pareto fronts or carrying out the robustness analysis does not become excessively high.

We have illustrated our framework using two example case studies: epidemic control and cancer screening program design. Both case studies demonstrated the potential of our framework for generating insights into complex decision problems that would have been hard to obtain without the use of optimization and the ability to ultimately focus on a small set of representative solutions. Although the epidemic case study is more stylized by nature, it nonetheless allowed us to get qualitative insights into optimal solution pathways. For instance, the framework helped discover a so-called 'lockdown by testing' effect related to strategies that use mass testing with imperfectly accurate tests. Moreover, it showed that given adequate testing capacity, mass testing even with imperfect tests can dominate pure lockdown strategies as well as strategies based on perfect tests but lower capacities. On the other hand, in the cancer screening case study, the framework suggested much lower cut-off levels for positive FIT results than those currently used in the screening program.

The main limitation of our proposed framework is the high computational effort required to solve the sets of Pareto optimal solutions and to perform the robustness analysis in a situation where the dynamics of the underlying processes and decision variables are captured by a complex, nonlinear model. In many cases long yet manageable computation times (such as the fourday computation time in the epidemic control case study) do not constitute a major barrier for the application of our framework given that it is intended to support large-scale and infrequent (if not one-off) decisions. If needed, this effort could be decreased by reducing precision in the optimization, or even in the process model. The complexity of the system description is a balancing act between the numerical tractability of the optimization algorithm and the realism of the process model. Often using even inexact algorithms or heuristics to solve an optimization problem that correctly represents the decision at hand can yield considerable benefits compared to using exact algorithms to solve a much simpler problem, or not using an optimization approach at all. Therefore, simplifying the process model should be carefully considered because this would inevitably reduce the real-life relevance of the model results. Finally, the visualization of solutions and objectives is an important phase in the proposed framework. Yet, such visualizations in more than three dimensions are known to be difficult. However, there exist techniques, such as animations and projections, that can help in this task.

This research opens up several interesting avenues for future work. First, it would be important to test the framework in different kinds of contexts in collaboration with real DMs. Potential contexts include, for instance, environmental decision-making and energy policy decisions. These kinds of context-specific applications are likely to reveal limitations in the applicability of our framework as it is currently presented. Some of these limitations could be overcome by improving the ways in which the stages in our framework are carried out. For instance, compared to how we conducted the robustness-based pruning of the Pareto fronts in our case studies, more advanced sampling techniques could be required in reallife applications to enable a larger sample size. Moreover, sensitivity analyses could be extended to cover not only the process model parameters but also the decision profiles. In the context of the COVID-19 case study, for instance, these kinds of sensitivity analyses could be used to examine the impact of small changes in the timing and strength of controls on relevant risk metrics.

## **Declaration of Competing Interest**

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Acknowledgements

Part of the research was developed in the Young Scientists Summer Program at the International Institute for Applied Systems Analysis, Laxenburg (Austria) with financial support from the Academy of Finland (grant number 336641). This research has also been supported by The Foundation for Economic Education (grant number 16-9442) and The Paulo Foundation. The optimization and sensitivity calculations were performed using computer resources within the Aalto University School of Science "Science-IT" project.

#### Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.ejor.2022.09.036.

#### References

- Aleta, A., Martin-Corral, D., y Piontti, A. P., Ajelli, M., Litvinova, M., Chinazzi, M., ... Merler, S., et al., (2020). Modelling the impact of testing, contact tracing and household quarantine on second waves of COVID-19. *Nature Human Behaviour*, 4(9), 964–971.
- Araz, O. M., Lant, T., Fowler, J. W., & Jehn, M. (2013). Simulation modeling for pandemic decision making: A case study with bi-criteria analysis on school closures. *Decision Support Systems*, 55(2), 564–575.
- Basurto, A., Dawid, H., Harting, P., Hepp, J., & Kohlweyer, D. (2020). Economic and epidemic implications of virus containment policies: insights from agent-based simulations. Working paper, *Bielefeld Working Papers in Economics and Management*.
- Berger, D., Herkenhoff, K., & Mongey, S. (2020). An SEIR infectious disease model with testing and conditional quarantine. Technical report, *Federal Reserve Bank* of Minneapolis.
- Caulkins, J., Grass, D., Feichtinger, G., Hartl, R., Kort, P. M., Prskawetz, A., ... Wrzaczek, S. (2020). How long should the covid-19 lockdown continue? *PLoS ONE*, 15(12), e0243413.
- Caulkins, J. P., Grass, D., Feichtinger, G., Hartl, R. F., Kort, P. M., Prskawetz, A., ... Wrzaczek, S. (2021). The optimal lockdown intensity for covid-19. *Journal of Mathematical Economics*, 93, 102489.
- Chaturvedi, A., Green, P. E., & Caroll, J. D. (2001). K-modes clustering. Journal of classification, 18(1), 35–55.
- da Cruz, A. R., Cardoso, R. T., & Takahashi, R. H. (2011). Multiobjective dynamic optimization of vaccination campaigns using convex quadratic approximation local search. In International conference on evolutionary multi-criterion optimization (pp. 404–417). Springer.
- Da Cruz, A. R., Cardoso, R. T., & Takahashi, R. H. (2009). Multi-objective design with a stochastic validation of vaccination campaigns. *IFAC Proceedings Volumes*, 42(2), 289–294.
- Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2), 182–197.
- Dellnitz, M., & Witting, K. (2009). Computation of robust pareto points. International Journal of Computing Science and Mathematics, 2(3), 243–266.
- DeNegre, A. A., Myers, K., & Fefferman, N. H. (2020). Impact of chemorophylaxis policy for aids-immunocompromised patients on emergence of bacterial resistance. *PLoS ONE*, 15(1), e0225861.
- Eichfelder, G., Krger, C., Schbel, A., & Eichfelder, G. (2015). Multi-objective regularization robustness.
- Falke, T., Krengel, S., Meinerzhagen, A.-K., & Schnettler, A. (2016). Multi-objective optimization and simulation model for the design of distributed energy systems. *Applied Energy*, 184, 1508–1516.

- Finnish colorectal cancer screening expert groups (2021). Finnish colorectal cancer screening protocol. https://syoparekisteri.fi/assets/files/2021/11/Protocol-forand-tests-used-in-colorectal-cancer-screening.pdf.
- Franco, L. A., & Montibeller, G. (2010). Facilitated modelling in operational research. European Journal of Operational Research, 205(3), 489–500.
- Friedrich, T., Kroeger, T., & Neumann, F. (2011). Weighted preferences in evolutionary multi-objective optimization. In Australasian joint conference on artificial intelligence (pp. 291–300). Springer.
- Groetzner, P., & Werner, R. (2021). Multiobjective optimization under uncertainty: A multiobjective robust (relative) regret approach. European Journal of Operational Research, forthcoming.
- Hoertel, N., Blachier, M., Blanco, C., Olfson, M., Massetti, M., Rico, M. S., ... Leleu, H. (2020). A stochastic agent-based model of the SARS-CoV-2 epidemic in France. *Nature Medicine*, 26(9), 1417–1421.
  Holzmann, T., & Smith, J. C. (2018). Solving discrete multi-objective optimization
- Holzmann, T., & Smith, J. C. (2018). Solving discrete multi-objective optimization problems using modified augmented weighted Tchebychev scalarizations. *European Journal of Operational Research*, 271(2), 436–449. https://doi.org/10.1016/j. ejor.2018.05.036.
- Howard, R. A., & Matheson, J. E. (2005). Influence diagrams. *Decision Analysis*, 2(3), 127–143.
- IHME COVID-19 forecasting team (2021). Modeling COVID-19 scenarios for the United States. *Nature Medicine*, 27(1), 94.
- Kasprzyk, J. R., Nataraj, S., Reed, P. M., & Lempert, R. J. (2013). Many objective robust decision making for complex environmental systems undergoing change. *Environmental Modelling & Software*, 42, 55–71.
- Kaufmann, L., & Rousseeuw, P. (1987). Clustering by means of medoids. In Statistical data analysis based on the l1 norm and related methods (pp. 405–416).
- Klein, R. W., Dittus, R. S., Roberts, S. D., & Wilson, J. R. (1993). Simulation modeling and health-care decision making. *Medical Decision Making*, 13(4), 347–354.
- Lempert, R. J., Groves, D. G., Popper, S. W., & Bankes, S. C. (2006). A general, analytic method for generating robust strategies and narrative scenarios. *Management Science*, 52(4), 514–528.
- Li, Z., Liao, H., & Coit, D. W. (2009). A two-stage approach for multi-objective decision making with applications to system reliability optimization. *Reliability En*gineering & System Safety, 94(10), 1585–1592.
- Lin, J., Liu, M., Hao, J., & Jiang, S. (2016). A multi-objective optimization approach for integrated production planning under interval uncertainties in the steel industry. *Computers & Operations Research*, 72, 189–203.
- Meng, K., Lou, P., Peng, X., & Prybutok, V. (2017). Multi-objective optimization decision-making of quality dependent product recovery for sustainability. *International Journal of Production Economics*, 188, 72–85.
- Miettinen, K., & Mäkelä, M. (1995). Interactive bundle-based method for nondifferentiable multiobjeective optimization: Nimbus. Optimization, 34(3), 231–246.

- Miller, R., Ewy, W., Corrigan, B. W., Ouellet, D., Hermann, D., Kowalski, K. G., ... El-Kattan, A., et al., (2005). How modeling and simulation have enhanced decision making in new drug development. *Journal of Pharmacokinetics and Pharmacodynamics*, 32(2), 185–197.
- Neuvonen, L., Dillon, M., Vilkkumaa, E., Salo, A., Jäntti, M., & Heinävaara, S. (2022). Decision programming for optimizing the finnish colorectal cancer screening program. arXiv preprint arXiv:2206.13830doi:10.48550/ARXIV.2206.13830.
- Petchrompo, S., Wannakrairot, A., & Parlikad, A. K. (2021). Pruning Pareto optimal solutions for multi-objective portfolio asset management. *European Journal of Operational Research, forthcoming.* Rangaiah, G. P. (2016). Multi-objective optimization: Techniques and applications in
- Rangaiah, G. P. (2016). Multi-objective optimization: Techniques and applications in chemical engineering: vol. 5. World Scientific.
- Salo, A., & Hämäläinen, R. P. (2010). Preference programming-multicriteria weighting models under incomplete information. In *Handbook of multicriteria analysis* (pp. 167–187). Springer.
- Salo, A. A., & Hämäläinen, R. P. (1995). Preference programming through approximate ratio comparisons. European Journal of Operational Research, 82(3), 458–475.
- Schöbel, A., & Zhou-Kangas, Y. (2021). The price of multiobjective robustness: Analyzing solution sets to uncertain multiobjective problems. *European Journal of Operational Research*, 291(2), 782–793.
- Sudeng, S., & Wattanapongsakorn, N. (2015). Post Pareto-optimal pruning algorithm for multiple objective optimization using specific extended angle dominance. *Engineering Applications of Artificial Intelligence*, 38, 221–236.
- Sudeng, S., & Wattanapongsakorn, N. (2016). A knee-based multi-objective evolutionary algorithm: an extension to network system optimization design problem. *Cluster Computing*, 19(1), 411–425.
- Taboada, H. A., & Coit, D. W. (2007). Data clustering of solutions for multiple objective system reliability optimization problems. Quality Technology & Quantitative Management, 4(2), 191–210.
- Van Der Zee, D.-J., & Van Der Vorst, J. G. (2005). A modeling framework for supply chain simulation: Opportunities for improved decision making. *Decision Sciences*, 36(1), 65–95.
- de Vos, N. J. (2015–2021). kmodes categorical clustering library. https://github.com/ nicodv/kmodes.
- Wismans, L. J., Brands, T., Van Berkum, E. C., & Bliemer, M. C. (2014). Pruning and ranking the Pareto optimal set, application for the dynamic multi-objective network design problem. *Journal of Advanced Transportation*, 48(6), 588–607.
- Yu, S., Zheng, S., Gao, S., & Yang, J. (2017). A multi-objective decision model for investment in energy savings and emission reductions in coal mining. *European Journal of Operational Research*, 260(1), 335–347.
- Zio, E., & Bazzo, R. (2011). A clustering procedure for reducing the number of representative solutions in the Pareto front of multiobjective optimization problems. *European Journal of Operational Research*, 210(3), 624–634.