VALUE IN HEALTH 21 (2018) 1019–1028

Available online at www.sciencedirect.com

ScienceDirect

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

ISPOR Report

Application of Constrained Optimization Methods in Health Services Research: Report 2 of the ISPOR Optimization Methods Emerging Good Practices Task Force

William Crown, PhD[1], Nasuh Buyukkaramikli, PhD[2], Mustafa Y. Sir, PhD[3], Praveen Thokala, PhD[4], Alec Morton, PhD[5], Deborah A. Marshall, PhD[6,7,8], Jonathan C. Tosh, PhD[9], Maarten J. Ijzerman, PhD[10,11], William V. Padula, PhD, MSc[12], Kalyan S. Pasupathy, PhD[3,*]

[1]OptumLabs™, Boston, MA, USA; [2]Institute of Medical Technology Assessment (iMTA), Erasmus University Rotterdam, Rotterdam, The Netherlands; [3]Health Care Policy and Research, Information & Decision Engineering, Mayo Clinic Kern Center for the Science of Health Care Delivery, Rochester, MN, USA; [4]University of Sheffield, Sheffield, UK; [5]Department of Management Science, Strathclyde Business School, University of Strathclyde, Glasgow, Scotland, UK; [6]Health Services & Systems Research, University of Calgary, Calgary, Alberta, Canada; [7]Alberta Bone & Joint Health Institute, Department Community Health Sciences, Faculty of Sciences, Faculty of Medicine, University of Calgary, Calgary, Alberta, Canada; [8]Cumming School of Medicine, University of Calgary, Calgary, Alberta, Canada; [9]DRG Abacus, Manchester, UK; [10]University of Twente, Department Health Technology & Services Research, Enschede, The Netherlands; [11]Luxembourg Institute of Health, Health Economics and Evidence Synthesis Unit, Strassen, Luxembourg; [12]Department of Health Policy & Management, Johns Hopkins Bloomberg School of Public Health, Baltimore, MD, USA

ABSTRACT

Background: Constrained optimization methods are already widely used in health care to solve problems that represent traditional applications of operations research methods, such as choosing the optimal location for new facilities or making the most efficient use of operating room capacity. **Objectives:** In this paper we illustrate the potential utility of these methods for finding optimal solutions to problems in health care delivery and policy. To do so, we selected three award-winning papers in health care delivery or policy development, reflecting a range of optimization algorithms. Two of the three papers are reviewed using the ISPOR Constrained Optimization Good Practice Checklist, adapted from the framework presented in the initial Optimization Task Force Report. The first case study illustrates application of linear programming to determine the optimal mix of screening and vaccination strategies for the prevention of cervical cancer. The second case illustrates application of the Markov Decision Process to find the optimal strategy for treating type 2 diabetes patients for hypercholesterolemia using statins. The third paper (described in Appendix 1) is used as an educational tool. The goal is to describe the characteristics of a radiation therapy optimization problem and then invite the reader to formulate the mathematical model for solving it. This example is particularly interesting because it lends itself to a range of possible models, including linear, nonlinear, and mixed-integer programming formulations. From the case studies presented, we hope the reader will develop an appreciation for the wide range of problem types that can be addressed with constrained optimization methods, as well as the variety of methods available. **Conclusions:** Constrained optimization methods are informative in providing insights to decision makers about optimal target solutions and the magnitude of the loss of benefit or increased costs associated with the ultimate clinical decision or policy choice. Failing to identify a mathematically superior or optimal solution represents a missed opportunity to improve economic efficiency in the delivery of care and clinical outcomes for patients. The ISPOR Optimization Methods Emerging Good Practices Task Force's first report provided an introduction to constrained optimization methods to solve important clinical and health policy problems. This report also outlined the relationship of constrained optimization methods relative to traditional health economic modeling, graphically illustrated a simple formulation, and identified some of the major variants of constrained optimization models, such as linear programming, dynamic programming, integer programming, and stochastic programming. The second report illustrates the application of constrained optimization methods in health care decision making using three case studies. The studies focus on determining optimal screening and vaccination strategies for cervical cancer, optimal statin start times for diabetes, and an educational case to invite the reader to formulate radiation therapy optimization problems. These illustrate a wide range of problem types that can be addressed with constrained optimization methods.

Keywords: Health care delivery, health services, health policy, medical decision making, operations research, constraints, optimal, optimization.

* Address correspondence to: Kalyan S. Pasupathy, PhD, Healthcare Policy & Research, Information & Decision Engineering, Clinical Engineering Learning Labs, Mayo Clinic Kern Center for the Science of Health Care Delivery, Rochester, MN, USA. Tel.: +(507) 293-2512.
E-mail: pasupathyk@mayo.edu.

1098-3015$36.00 – see front matter Copyright © 2018, ISPOR–The Professional Society for Health Economics and Outcomes Research. Published by Elsevier Inc.
https://doi.org/10.1016/j.jval.2018.05.003

Background to the Task Force

The proposal to initiate an ISPOR Good Practices for Outcomes Research Task Force was evaluated by the ISPOR Health Science Policy Council then recommended to the ISPOR Board of Directors for approval.

The task force was comprised of international subject matter experts representing a diverse range of stakeholder perspectives (academia, research organizations, government, regulatory agencies and commercial entities). The task force met approximately every five weeks by teleconference and in person at ISPOR conferences. All task force members reviewed many drafts of the report and provided frequent feedback in both oral and written comments.

To ensure that ISPOR Good Practices Task Force Reports are consensus reports, findings and recommendations are presented and discussed at ISPOR conferences. In addition, the first and final draft reports are circulated to the task force's review group. All reviewer comments are considered. Comments are addressed as appropriate in subsequent versions of the report. Most are substantive and constructive improving the report.

Introduction

There are often many different options for improving health care policy or improving current practice in health care organizations. The optimal solution among those options, i.e., the solution that best achieves a defined goal, such as maximizing patient quality of life or minimizing patient waiting time for services, may not be readily apparent. Constrained optimization methods use mathematical techniques to help efficiently and systematically identify the best (optimal) of all possible solutions to a problem while considering the relevant constraints, such as budget limits or staffing capacity.

Of course, mathematically optimal solutions to all problems are not always feasible; other nonquantifiable criteria, such as political barriers that cannot be accounted for by defined constraints, have to be considered. However, optimization techniques can still be highly informative to decision makers in providing insights about optimal target solutions and the magnitude of the loss of benefit or increased costs associated with the ultimate policy choice. In health care, failing to identify a mathematically superior or optimal solution represents a missed opportunity to improve economic efficiency in the delivery of care and clinical outcomes for patients.

The ISPOR Optimization Methods Emerging Good Practices Task Force provided an introduction to constrained optimization methods to solve important health policy and clinical problems in its first report [1]. The previous report outlined the relationship of constrained optimization methods relative to traditional health economic modeling and simulation models and identified some of the major variants of constrained optimization models, such as linear programming, dynamic programming, integer programming, and stochastic programming.

In addition, the report graphically illustrated the formulation and solution of a straightforward integer program to maximize health benefit subject to a budget constraint. Further, it explained the steps in an optimization process: 1) structuring the problem; 2) formulating the mathematical model; 3) developing the model; 4) validating the model; 5) selecting the optimization method; 6) performing the optimization and conducting sensitivity analysis; 7) reporting results; and 8) using the results for decision making.

The principal objective of this second Optimization Task Force Report is to illustrate the application of constrained optimization methods in health care decision making. To identify relevant examples, we began by searching for award-winning health care papers from the Institute for Operations Research and Management Sciences (INFORMS) and the Association for European Operations Research Societies (EURO). From these papers, we then selected examples with models relevant for health economic policy or clinical decision making. Finally, we endeavored to select papers that collectively illustrated a variety of different constrained optimization methods. The three papers that received the most votes from the task force members were selected.

In this report, two of these three papers are compared with the steps in formulating, solving, validating, reporting, and using optimization models originally published as Table 3 in the first Optimization Emerging Good Practices Task Force Report. A slightly modified version of this previous table is presented as the ISPOR Constrained Optimization Good Practices Checklist (Table 1) in the current report. The first case study illustrates the application of linear programming to determine the optimal mix of screening and vaccination strategies for the prevention of cervical cancer [2].

Table 1 – ISPOR Constrained Optimization Task Force Good Practices Checklist

Stage	Step	Description
Modeling	Problem structuring	Specify the objective(s) and constraints, identify decision variables and parameters, and list and appraise model assumptions
	Mathematical formulation	Present the objective function(s) and constraints in mathematical notation using decision variables and parameters
	Model development	Program the model in software to estimate the objective function(s) and constraints using decision variables and parameters as inputs
	Model validation	Ensure the model is appropriate for evaluating different combinations of decision variables and parameters
Optimization	Select optimization method	Choose an appropriate optimization method and algorithm on the basis of characteristics of the problem.
	Perform optimization/ sensitivity analysis	Use the optimization algorithm to search for the optimal solution and examine the performance of the optimal solution for reasonable sets of parameters
	Report results	Report the results of the optimal solution and sensitivity analyses
	Decision making	Interpret the optimal solution and use it for decision making

Source: Crown et al. [1], Table 3, p. 315.

The second case illustrates application of the Markov Decision Process to find the optimal strategy for treating type 2 diabetes patients for hypercholesterolemia using statins [3].

Finally, the third paper (described in Appendix 1) is used as an educational tool. The goal is to describe the characteristics of a radiation therapy optimization problem and then invite the reader to formulate the mathematical model for solving it. This example is interesting because it lends itself to a range of possible models, including linear, nonlinear, and mixed-integer programming formulations. (Detailed formulations for each model are provided in Appendix 1.)

Although we are clearly limited in the number of permutations we can present with these three cases, we hope the reader will develop a sense of the wide range of problem types that can be addressed with constrained optimization methods as well as the variety of methods available.

Overview of Applications of Constrained Optimization in Healthcare

Constrained optimization methods are already widely used in health care areas, such as choosing the optimal location for new facilities, making the most efficient use of operating room capacity, workforce planning, etc. They can also be instrumental in guiding clinical decision making in actual clinical practice where health professionals and patients face constraints, such as proximity to treatment centers, health insurance benefit designs, and the limited availability of health resources.

Optimization is also beneficial for planning health care expenditure. An obvious example is the resource allocation problem faced by a planner with a number of investment opportunities, but a fixed budget inadequate to fund all available opportunities [4]. Perhaps the simplest case of this is where the investment opportunities are incremental to current care and fall into distinct categories (e.g., children's services, cardiovascular disease, cancer, respiratory disease, and mental health) with separate budgets [5]. In this situation, decisions about investments in different clinical areas can be made independently of one another.

However, more commonly the health care budget needs to be allocated across different conditions. The problem of choosing the best set of investment opportunities to fund under a fixed budget constraint in order to meet an objective, such as maximizing total quality-adjusted life-years (QALYs), can be addressed as an optimization problem [6]. Given a number of eligible interventions and a fixed budget, optimization can be used to solve resource allocation problems.

In fact, the task central to health economic analysis, i.e., evaluating whether the incremental cost-effectiveness ratio (ICER) of an intervention is below a critical threshold, can be shown to be related to budget constrained optimization. According to the theoretical definition, under a strict set of assumptions, the threshold represents the inverse of the shadow price of the budget constraint; the shadow price is defined as how much the objective (QALYs) would increase for a one-unit increase in the constraint (budget) [7].

Other resource allocation problems may be even more complicated. There may be significant and complex interactions between different investments; and there may be additional constraints to be considered, such as limits on the number of staff or bed capacity [8]. For example, consider the case of allocating resources for the prevention and cure of an infectious disease such as HIV, hepatitis C, tuberculosis, malaria, or polio [9,10]. If the planner invests in vaccination, there may be fewer cases to treat in the future, and so investment in highly capital-intensive treatment facilities may be wasted. On the other hand, vaccination is itself costly, and if the disease has a low prevalence, it may be more cost-effective to target the treatment [11]. Additional work in this area is being performed and will appear in a future issue of *Value in Health*.

Optimizing investment in such infectious disease programs is more complicated because they may involve making multiple runs of a state-of-the-art simulation [12,13] of the infectious disease dynamics to plot out how the particular patterns of resource allocation perform against the objective of minimizing the total number of cases or maximizing the probability of achieving disease eradication. For a review of mathematical approaches to infectious disease prediction and control, see Dimitrov and Meyers [14].

In other settings, the critical resource(s) might not be money; for example, when allocating donated organs such as kidneys, not every kidney will be compatible with every donor. In addition, the medical condition of the eligible recipients will be different; some will be more urgent than others. In this case, the underlying problem can be categorized as a matching problem [15]. In matching problems, not everyone will get the best match. However, the objective with kidney allocation is generally to ensure that as few people and kidneys as possible are left unmatched (i.e., patients without kidneys, kidneys without patients). Bertsimas et al. [16] present a discussion about how to incorporate fairness in such problems. Some measures of deservingness, e.g., time on waiting list, may be incorporated in the objective function. Nevertheless, some fairness considerations may also be included as constraints, e.g., at least x% of transplants should go to patients of a certain blood type. The 2012 Nobel Prize in Economics was awarded to Shapley and Roth, in part for their work in stable matchings applied to organ donation.

Other clinical problems where optimization can be applied relate to problems of disease management, e.g., timing of the initiation of treatment, or the sequence of treatments. The promise of health gain from treatment must be balanced against reasons for holding off treatment, which may include cost, undesirable side effects, and emergent drug resistance. It may be the case that there is an optimal stage in the disease prognosis or point in the patient's life cycle where the balance shifts from favoring nonintervention to favoring treatment. The Markov Decision Process (MDP) approach provides an ideal framework [17] to study such problems for identifying critical initiation points. This framework has been used to analyze timing decisions in diseases as diverse as HIV, diabetes, and breast cancer [3,18,19]. Optimization methods can be applied to identify the optimal treatment sequences when a large number of treatments are available; for example, in rheumatoid arthritis [20].

Finally, constrained optimization methods have also been applied to disease diagnosis [21,22], the development of optimal treatment algorithms [23,24], and the optimal design of clinical trials [25]. Health technology assessment using tools from constrained optimization methods is also gaining popularity [26]. We also encourage the readers to refer to the initial ISPOR Optimization Emerging Good Practices Task Force Report, which presented a more comprehensive overview of the different applications for which optimization techniques can be used [1].

Steps in an Optimization Process

Table 1 reproduces the steps of the optimization process previously presented in the initial Optimization Task Force Report. It is reproduced here as the ISPOR Constrained Optimization Good Practice Checklist to reduce reader burden because the two case studies and the educational example will all be discussed in light of this framework. The primary purpose of Table 1 is to support the design of optimization studies by prompting the user to

report and justify the choices made at each step of the process. It should be noted that the steps outlined in Table 1 do not need to be conducted sequentially. In fact, most of the optimization studies involve performing these steps in an *iterative* manner to solve the problem. Along with guiding the design of optimization studies, Table 1 can also be used to support the critique and quality assessment of published optimization studies. The steps in Table 1 are described in detail in the text below.

Problem Structuring

The first step is to determine if constrained optimization is an appropriate methodology to address the problem at hand. It involves identifying if there are any quantifiable constraints and whether a specific goal can be achieved by changing some decision variables. This problem structuring phase should be done in consultation with the key stakeholders and decision makers to ensure that the optimization problem is appropriately specified. This will ensure that the objective functions and constraints are appropriate and get their "buy-in" to change the decision variables in order to achieve an optimal solution. A clear textual description of the decision problem should be reported to and validated with the key stakeholders and decision makers.

Mathematical Formulation

This involves converting the textual description into a mathematical representation of the optimization problem. Objective function(s) and constraints need to be defined in analytical form as a function of decision variables and parameters. Note that decision variables are changed during optimization iterations in order to identify the optimal solution, although parameters remain fixed. The number and type of decision variables (continuous or discrete) as well as the parameters need to be justified. The type of objective function (single objective or multiobjective, linear or nonlinear, stochastic or deterministic) and the type of estimation (analytical estimation or simulation modeling for complex problems) need to be specified. Similarly, for constraints, the number of constraints and the type of estimation used for the constrained quantity need to be reported and justified. The sources and the values of the parameters used to estimate the objective function(s) and constraints also need to be justified. The mathematical representation of the optimization problem should be reported after validation with the key stakeholders and decision makers.

Model Development

This involves programming the model in software to estimate the objective function(s) and constraints, using decision variables and parameters as inputs. It should be noted that in some instances, the analytical form of the mathematical formulation can be programmed directly because the mathematical formulation sufficiently defines the relationships between objective function(s)/constraints and decision variables/parameters. However, in other instances, a simulation model needs to be developed to estimate the objective function(s)/constraints. Models should be designed so that the objective function can be evaluated based on the full range of possible decision variables (the feasible region or search space). The model structure and assumptions should be reported and validated with the key stakeholders and decision makers. The initial mathematical formulation and model development steps affect the specification of the particular optimization method to be applied. These steps are *closely related and interdependent*. This is one important reason why the steps in optimization do not always have to follow the order described in Table 1.

Model Validation

Before optimization is undertaken, the underlying model needs to be verified and validated to ensure the robustness of the results for different analyses performed. This means that the model should be consistent with reality within specified tolerances. Once the model has been developed to the point where it is producing estimates, the model code also needs to be checked to ensure the model results are valid. In the case of models that represent an analytical formulation directly, this is relatively straightforward because this involves checking the specific model results used as parameters for estimating the objective function and constraints.

However, when a simulation model is used to evaluate the objective function, this would necessitate a combined approach of simulation-optimization [27,28]. This is a bit more difficult because it involves checking the model results for all combinations of decision variables. Metamodeling techniques [29], i.e., modeling the simulation model outputs as functions of simulation inputs, can circumvent getting the simulation results for all variables in the parameter space. These topics are beyond the scope of this report; we suggest reviewing Sargent [30] and Law [31].

Modelers are encouraged to validate the model results in different parts of the decision variable space to have enough confidence that the model used is appropriate for optimization [32,33]. This should also involve asking the key stakeholders and decision makers to check the model results for face validity.

Select Optimization Method

The choice of optimization method needs to be justified on the basis of the type of decision variables (continuous or discrete), the type of objective function (single objective or multi-objective, linear or nonlinear, stochastic or deterministic), and the type of constraints (single vs. multiple). The optimization algorithm or tool used also needs to be justified on the basis of the optimization method as well as the estimation type (analytical formulation vs. simulation optimization) and other relevant characteristics of the model (number of decision variables or transferability of the problem to other well-known problem types). The methods and tools chosen for optimization need to be reported and justified.

Perform Optimization/Sensitivity Analysis

This involves running the optimization model, identifying the optimal solution, and understanding the impact of alternative parameters on the optimal solution using sensitivity analyses. Settings used for the optimization, such as the convergence level required or the maximum number of iterations, need to be justified. In some problems, searching for the optimal solution might be computationally feasible, whereas in others, solving time increases to such an extent that the use of heuristics is justified.

As with decision modeling, optimization can have stochastic uncertainty in parameters and model structure. Stochastic optimization [34], robust optimization techniques [35], and sensitivity analyses can be used to deal with parameter uncertainty. However, structural uncertainty needs to be dealt with by thinking about the choices throughout the optimization process. For example, is a linear program really appropriate? Are the simplifications and assumptions appropriate? Further, to what extent is there a risk of a wrong or suboptimal decision being reached? The choice of decision variables, parameters, constraints, and model assumptions also needs to be structurally evaluated.

The optimal solution needs to be checked to identify if it is feasible and, if so, sensitivity analyses should be conducted. The optimization settings and the sensitivity analyses need to be explained to the key stakeholders or decision makers and reported in detail.

Report Results

This involves specifying the values of the decision variables, objective function, and constraints at the optimal solution for the base case analyses as well as the sensitivity analyses. The optimization results (i.e., optimal solution for the base case and sensitivity analyses) need to be reported and validated with the key stakeholders or decision makers. Also, the performance of the optimization tool or method, such as the time taken to identify the optimal solution, number of iterations required, and the convergence level (if applicable), needs to be reported. These results should be reported in a manner that is understandable and interpretable by relevant stakeholders and decision makers.

Decision Making

The meaning of the optimal solution should be explained to the decision makers. This involves converting the mathematical optimal solution into clear, concise plans for implementation. At this stage, the choices made at all the stages of modeling and optimization (i.e., type of model, data, and assumptions as well as design, settings, and others) need to be validated to ensure the results of the optimization problem are plausible and consistent with decision maker objectives. Also, the possibility of amending the decision variables to the values specified by the optimization process needs to be checked with the stakeholders to ensure that the implementation is feasible.

To reiterate, the results of the optimization should not be used mechanistically. It is the decision makers that implement the findings; hence, they should be comfortable with the methodology, data, and assumptions involved in the whole optimization process.

Optimization Case Studies

In this section, we consider two constrained optimization studies and compare their structure to the steps outlined in Table 1. The first case study focuses on resource allocation for the prevention and cure of infectious diseases, and the second illustrates the use of constrained optimization to guide optimal treatment initiation. These cases illustrate different modeling techniques as well as extensions of the application of constrained optimization methods beyond the typical realm of scheduling, shipping cost minimization, maximization of facility capacity, etc. Please note that the educational case study and the model formulations appear in Appendixes 1 and 2, respectively.

Case Study 1: Selecting a Mix of Prevention Strategies Against Cervical Cancer

Problem Structuring

Cervical cancer is the second most common cancer in women under 35 years old in the United Kingdom (UK). The objective of this study was to identify the optimal mix of primary and secondary prevention strategies for cervical cancer that achieves maximum reduction in cancer cases under budget and logistic constraints [2]. The authors applied the optimization model in two countries (UK and Brazil) with different health care organizations, epidemiology, screening practices, resource settings, and treatment costs. They considered two cervical cancer prevention strategies against human papillomavirus (HPV):

1. Primary prevention: Because an HPV infection is the most common cause of cervical cancer, HPV vaccination is a primary prevention strategy. Two HPV vaccines are currently available. Both vaccines have an efficacy of approximately 98% against the cervical cancer vaccine HPV types (HPV 16 and 18), but with a different cross-protection profile against oncogenic nonvaccine HPV types. The implementation of vaccination varies widely among countries with regard to the strategy selection (national immunization program or individual-based), the logistics (via a separately established vaccination setting or via primary health care), the age group targeted, and the gender selection (female only or all patients).
2. Secondary prevention: Cytology-based screening programs have contributed to a decrease of up to 80% in the incidence and mortality of cervical cancer in countries with a well-established, organized screening program. However, despite their potential, cytology-based screening programs sometimes have a limited impact due to factors such as sensitivity of the screening method (ability of the test to correctly identify those patients with the disease), treatment failure, and the level of resources required for an adequate follow-up of patients.

Four prevention strategies were evaluated: screening, vaccination, screening plus vaccination, and no prevention because these were the options available for cervical cancer prevention in the UK and Brazil at the time of the study. Only cytology-based screening was included in the model, with sensitivity estimates based on published literature. Different screening interval scenarios were explored, from every year to every 25 years. Women were screened only twice over their lifetime, with a 1-year increment between each scenario.

It was assumed that vaccination was administered at age 12 and induced lifelong protection against HPV. In total, 52 different strategies were tested for each country. These 52 strategies defined the full range of possible combinations of vaccination (not available or available) and screening interventions (not available or available, with intervals between screenings estimated from 1 year to 25 years in 1-year increments). The final scenarios can be listed as follows: (*scenario 1*: no screening & no vaccine; *scenario 2*: 1-year screening interval & no vaccine; *scenario 3*: 2-year screening & no vaccine; ..., *scenario 26*: 25-year screening & no vaccine; *scenario 27:* no screening & vaccine; *scenario 28:* 1-year screening & vaccine; *scenario 29:* 2-year screening & vaccine; ..., *scenario 52*: 25-year screening & vaccine).

Mathematical Formulation

The optimization model used a linear programming formulation consisting of a single linear objective function and multiple linear constraints. The model was continuous, allowing fractional values for the decision variables. It was static, meaning that the problem was solved once at steady state. Finally, the model was deterministic, which assumed that all the outputs were known and there was no stochastic variation.

Fifty-two decision variables, x_i, each representing the proportion of the population addressed by each strategy considered, $i = 1,2,..,52$, were used with separate identifiers for strategies involving screening and strategies involving vaccination in order to deal with screening and vaccination coverage constraints. Given the aim was to minimize the number of cervical cancer cases, the objective function was represented as the sum of the cervical cancer cases (at steady state for 100,000 women) for each strategy,CC_i, multiplied by the proportion of population receiving each strategy,x_i

The linear programming formulation for the cervical cancer prevention strategy optimization is given as the following:

$$\sum_{i=1}^{52} CC_i x_i \tag{1}$$

$$\sum_{i=1}^{52} b_i x_i \le B \quad (2)$$

(budget constraint)

$$0 \le x_i \le 1, \qquad for\ i=1,2,\ldots,52 \quad (3)$$

(strategy coverage bounds)

$$\sum_{i=1}^{52} x_i = 1 \quad (4)$$

(complete population distribution)

$$\sum_{i=2}^{26} x_i + \sum_{i=28}^{52} x_i \le Cov_1 \quad (5)$$

(screening coverage upper bound)

$$\sum_{i=27}^{52} x_i \le Cov_2 \quad (6)$$

(vaccination coverage upper bound)

$$x_1 \le \min(1-Cov_1, 1-Cov_2) \quad (7)$$

(upper bound on population with no coverage)

$$x_i \in R, for \qquad i=1,2,\ldots,52 \quad (8)$$

The model has five constraints: budget, strategy coverage, total population, screening, and vaccination coverage limits. The first constraint is to ensure that the sum of the cost for each strategy (at steady state for 100,000 women), b_i, multiplied by the proportion of the population receiving each strategy x_i, is less than the overall budget constraint, B. The strategy coverage constraint ensures that the proportion of each strategy is between zero and one. The complete population distribution constraint guarantees that all 52 variables add up to 1. That is, the sum of the proportion of the population receiving each strategy should reflect the entire population.

Also, the sum of the proportion of the population receiving strategies including screening should be less than the government-imposed screening coverage limit, Cov_1,. Similarly, the sum of the proportion of the population receiving strategies including vaccination should be less than the externally (e.g., government) imposed vaccination coverage limit,Cov_2. Note that the parameters CC_i and b_i are derived from the Markov cohort model (see details below) for each strategy i.

Model Development

The mathematical formulation described above used the outputs of a health economic Markov cohort model (number of cervical cancer cases [CC_i] and total costs [b_i] for each strategy [i]) as input parameters. The Markov cohort model describes the population level's natural history of cervical cancer for the evaluation of the clinical and economic consequences of different prevention strategies. The model considers a population of 100,000 women under a given prevention strategy at steady state level. The Markov model consists of the following states: no HPV infection, HPV infection, cervical intraepithelial neoplasia (CIN) stages, cancer, and death (both cancer and noncancer related).

Once patients are infected with HPV, individuals can progress and regress from HPV infection and CIN stages. Vaccination is assumed to reduce the HPV infection rates, and detection through screening provides the possibility of the treatment of CIN. Overall vaccine efficacy in the UK and Brazil was calculated from the country-specific proportions of each HPV type in cervical cancer. Other clinical and cost inputs were specified for each of these two countries.

The time horizon of the optimization problem was 1 year, and both the health and cost outcomes across the whole population were derived from the lifetime cohort results from the Markov model.

The model was run separately (for both countries) with a cohort of women over their lifetime for each one of the 52 scenarios described above. The results of each scenario were used to estimate the number of cervical cancer cases and total costs expected over 1 year at steady state for 100,000 women. The estimated number of cervical cancer cases (CC_i) and total costs (b_i) of each of the 52 prevention strategies were then used as input parameters for the optimization model.

Model Validation

No validation effort was reported for either the health economic model or the optimization model.

Select Optimization Method

Due to the relatively small size of the linear programming formulation described above (a total of 52 decision variables and 57 constraints), a standard primal simplex method was chosen to solve the problem.

Perform Optimization/Sensitivity Analysis

This optimization problem was programmed in Microsoft Excel as a linear program and solved using the Solver Add-in. This tool uses the simplex method to identify the optimal mix of the 52 cervical cancer prevention strategies to minimize the expected cervical cancer cases under a fixed budget as well as screening and vaccination coverage constraints. The optimization model was solved twice using separate parameter sets, reflecting the settings in UK and Brazil.

The base-case analysis assumed that the maximum screening coverage is the prevaccination coverage rate (65% in the UK and 50% in Brazil), maximum vaccination coverage was set to 80%, and the overall budget was the prevaccination budget allotted to screening and treatment of cervical lesions. No explanation was given as to why these maximum coverage rates were chosen in the base-case.

Sensitivity analyses were performed to understand the effect of altering the budget or the achievable screening or vaccination coverages (the constraints in the model) as well as the duration of vaccine protection, which was one of the parameters in the economic modeling. The budget constraint was varied from a 20% reduction to a 150% increase over the prevaccination levels, and the screening and the vaccination coverage levels were varied from 0% to 100%.

Report Results

The optimal mix of strategies in the UK was 65% vaccination plus screening, with a screening interval of 6 years, and 15% vaccination alone. In Brazil, the optimal mix was 50% vaccination plus screening, with a screening interval of 5 years, and 30% vaccination alone. These optimal mixes of strategies would result in a reduction of cervical cancer by 41% in the UK and 54% in Brazil from prevaccination levels with no budget increase. It can be easily observed that in both countries, the optimal coverage rates for both preventive interventions are at the maximum levels permitted in the model.

In the sensitivity analyses, increasing the budget permits a shortening of the screening interval, but the effect on the reduction in cervical cancer cases is modest and tends to reach an early plateau. Vaccination alone (screening coverage set to 0%) could provide a reduction in cervical cancer cases compared with the prevaccination situation of screening alone with a lower budget. In both countries, the effect of reduced vaccine efficacy

duration (25 years compared with lifetime) still results in a reduction in cervical cancer compared with the prevaccination strategy, but not as much as the base-case analysis. In both countries, a sharp reduction in the expected number of cervical cancers is seen when the vaccine coverage rate exceeds the maximum screening coverage rate or when screening coverage rate exceeds the maximum vaccine coverage rate while maintaining the budget (treatment and prevention) constraint.

Decision Making

In this case study, within the same budget, results of the optimization program show that it would be possible to substantially reduce the number of cervical cancer cases by implementing an optimal combination of HPV vaccination (80% coverage) and screening at prevaccination coverage (65% UK, 50% Brazil) while extending the screening interval to every 6 years in the UK and 5 years in Brazil.

Optimization models can be used to determine the optimal mix of primary and secondary prevention strategies minimizing cervical cancer burden under budget and logistic and infrastructure constraints. The key strength of optimization modeling is its ability to evaluate multiple combinations of different interventions and identify the mix that provides the maximum expected health benefit (reduction in cervical cancer cases) at the expected costs within the available budget. In addition, it allows the decision maker to set constraints reflecting local conditions, such as a limited available budget or limited achievable coverage rates.

In this paper, the optimization model uses the health economic model outcomes as its input parameters. Therefore, the validity of the optimization results is based on the validity of the health economic model. Furthermore, the implementation issues, such as how it will be decided who will receive vaccination, screening, or both, were not discussed. In its current form, the optimization model is used more to demonstrate the potential value of adding vaccination strategy and to coordinate this addition with the existing screening practices in the UK and Brazilian health systems.

The epidemiologic model developed to estimate the outcomes of the HCV vaccination program was a static cohort model. This allowed using the simpler linear programming approach to solve the optimization problem but might have underestimated the impact of vaccination in reducing disease cases among those unvaccinated. A dynamic transmission model could have accounted for these effects, however this might have complicated the optimization process and might risk the model to be less transparent to decision makers. Methods need to be developed to allow consideration of secondary transmission while maintaining model transparency and increasing the usefulness of the results provided in the paper.

Case Study 2: Optimizing Statin Treatment Initiation Using MDP

Problem Structuring

Type 2 diabetes (T2D) leads to many chronic outcomes, including stroke, coronary heart disease (CHD), and kidney failure. This study focuses on the selection of T2D patients for statin therapy of hypercholesterolemia [3]. The market for statins is significant and remains burdensome to health system costs despite the availability of generics. Furthermore, there are a number of studies that report overprescribing (prescribing statins to those patients who only achieve marginal benefit) and underprescribing (not prescribing statins to those patients most likely to benefit). Given this debate, the aim of the study by Denton et al.was to identify the optimal time to initiate statin treatment for hypercholesterolemia in T2D patients [3].

The problem is set up using a MDP framework. Traditional health services research methods focus on efficacy or cost effectiveness at a snapshot in time to inform decisions, whereas MDP provides an in-depth modeling and understanding of optimal decisions at multiple time points over a patient's disease history. Due to the nature of the modeling, it provides the ability to personalize decisions, as opposed to one-size-fits-all policies and guidelines established for medical decisions. However, similar to other approaches, MDPs have assumptions based on data and/or the structure of the model.

Mathematical Formulation

The model optimizes a cost-reward function over time using an MDP. We recognize that MDPs are not commonly associated with constrained optimization because they typically do not have "constraints" in the same sense that the term is used in the mathematical programming literature (for example, in the previous case study). However, the ability of dynamic programming models to identify the optimal solution to the MDP (i.e., the optimal pattern of statin therapy initiation over time) provides an excellent example of a clinical use case for constrained optimization as long as one recognizes that constraints in an MDP are implicitly defined based on allowable transitions between states and/or available decisions within each state.

The structure of the model reflects shared decision making by providers and patients over time as a function of patient age, patient clinical history, and several health states. History is dependent on CHD or stroke as well as nine cholesterol levels pertaining to low, medium, or high high-density lipoprotein (HDL) and low-density lipoprotein (LDL) levels. Patient information aligning with the data across three major heart studies provides much higher sensitivity to the proper time to initiate and maintain a statin regimen. The MDP model determines at each epoch the optimal decision to maximize the overall rewards $v(s_t)$ while accounting for costs of all future states.

$$\text{Reward function:}\quad Max\, v(s_t) = E_{\underset{s}{\rightarrow}}\left[\sum_{t=1}^{T}\left(\lambda^{N^D t}\right) r(s_t, a(s_t))\right] \forall s_t \in \rightarrow S$$

where t is a time index for discrete decision epochs, s_t is an index for states at time period $t = 1,\ldots T$, $a\ (s_t)$ is the statin treatment decision at time $t = 1,\ldots T$, $\lambda \epsilon [0,1]$ discounts the objective function depicting reduced value of rewards in future years, and N^D is the number of years in a decision epoch.

Reward function for each time period: $r(s_t, a(s_t)) = N^D\left[R(s_t) - \left(CF^S(s_t) + CF^{CHD}(s_t)\right) - a(s_t)C^{ST}\right] - \left[C^S(s_t, a(s_t)) + C^{CHD}(s_t, a(s_t))\right]$

where N^D reflects the number of years in a decision epoch, $R\ (s_t)$ is the monetary value of quality life years, $C\ (St)$ is the annual cost of statin treatment in period t, $CF^s\ (s_t)$ is the annual follow-up care cost of stroke in period t, $CF^{CHD}\ (s_t)$ is the annual follow-up care cost of CHD event in period t, $C^S(s_T)$ is the one time cost of stroke occurring in period t.

Reward function for final time period: $r(s_T, a(s_T)) = N^D[R(s_T) - \left(CF^S(s_T) + CF^{CHD}(s_T)\right) - a(s_T)C^{ST}] - \left[C^S(s_T, a(s_T)) + C^{CHD}(s_T, a(s_T))\right] + E[PDHR \mid s_T, a(s_T)]$,

where $E[PDHR \mid s_T, a(s_T)]$ is the postdecision horizon expected reward. The authors separate the time horizon into a decision horizon and a postdecision horizon. While the decisions are only made during the decision horizon, the rewards from the postdecision horizon still need to be included. For instance, while the decision to initiate statin therapy is only until age 80, the rewards of treatment after age 80 need to be included in the model in Figure 2 in the article [3].

Model Development

The starting age of the patients in the model was 40, and it was assumed that the patients could start statin treatment at any point between 40 and 80 in 2-year increments. If all these treatment options were modeled as separate scenarios, as is common in both clinical trials and economic evaluations, the problem would soon become quite complicated, especially if these treatment options were compared incrementally.

However, using optimization techniques, one can identify a "single" optimal age for initiating statin treatment that maximizes the reward function. The authors interpreted reward in terms of expected net monetary benefit E(NMB) as a function of QALYs and cost and willingness-to-pay threshold (λ), that is:

$$E(NMB) = \Delta QALYs^{*}\lambda - \Delta Cost \quad (9)$$

Model Validation

The authors do not describe the model validation process, although it is clear from the manuscript's acknowledgments section that the authors interacted extensively with experts within the clinical system where the research was conducted as well as with external reviewers.

Select Optimization Method

The problem is set up using an MDP. The MDP framework is intended for dynamic streams of decisions (i.e., decisions made over time). The time horizon and the time steps are identified as indices for decision epochs. Each decision in the stream guides the evolution of the system being modeled (typically the patient's health in medical applications) and may enable or foreclose further decisions. The patient's health at each time point is typically the state, and the decisions or actions are identified. MDPs can be considered a hybrid between a Markov model and a decision tree.

Just as in Markov models, in an MDP, a patient's health state changes over time, transitioning from one discrete state to another according to a specified matrix of probabilities. However, typically in a Markov model, the decision maker has a choice between two or more treatment regimens to start the patient on initially. By contrast, in an MDP, the decision maker can make a choice about treatment in **every** time period. Thus, it is possible to model at a more granular level. At each time point, one may decide to start, stop, or switch treatments for as long as the patient survives. The constraints may involve the changes in states and/or the decisions. The transition from one state to another is characterized probabilistically.

In this study, the critical decision is when to start statins. Starting statins is taken to be a one-time irreversible decision. Thus, in each time period from age 40 to death—or age 80—there is a binary "start" or "delay" decision. Much of the complexity of the model is in the modeling of the health states. There are 324 health states describing various combinations of cholesterol and high-density lipoprotein levels (three each) as well as stroke and CHD states (six each). Transition probabilities are parameterized based on a proprietary clinical database. The objective function is a combination of health sector costs—such as the cost of treatment transacted between the provider, patient, and payer—and net monetary benefit, appropriately discounted over time. The risk of adverse events is modeled for comparison through three third-party risk models.

Different risk-prediction models have estimated probabilities of T2D complications in patients based on sociodemographic and environmental risk factors. These predictive models can specify the type of treatment to reduce the risk of comorbidity. The most common validated risk models from several large studies are the United Kingdom Prospective Diabetes Study (UKPDS), the Framingham Heart Study (United States [US]), and Archimedes, based on data trial results from the Heart Protection Study of 2002.

In particular, the article "Optimizing the Start Time of Statin Therapy for Patients with Diabetes" [3] aimed to identify the optimal decisions for individual patients based on their attributes, including age, gender, total cholesterol, and HDL. The authors also performed the analyses using the predictions from each of the three risk models above. Because the choice of the risk model may impact the treatment decision, they noted that the predictions from the models could be different.

Perform Optimization/Sensitivity Analysis

The solution method is based on a backward induction approach starting with the last epoch T. Knowing the optimal future actions, the optimal decision at the current epoch can be established using recursive optimality in the following equation:

Recursive optimality:

$$v(s_t) = \lambda^{N^D} \max[r(s_t, a(s_t)) + \sum_{\forall s_{t+1}} p(s_{t+1} | s_t, a(s_t)) v(s_{t+1})]$$

where

$p(s_{t+1} | s_t, a(s_t))$ is the state transition probability at time t given state s_t and action $a(s_t)$

Decision variable: $a(s_t) = \begin{cases} 1 \text{ if statin treatment is initiated} \\ 0 \text{ if statin treatment is delayed} \end{cases}$

where if $a_{t'} = 1$, *then* $a(s_t) = 1, \forall t > t'$

Where uncertainty in the model existed based on recommended statin starting therapy, the results of the optimization approach were tested for the low, medium, and high cost of statins across a willingness-to-pay threshold ranging from $25,000/QALY to $100,000/QALY in $25,000 increments. This additional analysis provides insight into the value of the model recommendations and whether the recommendation results from using a low- or high-value proposition as a starting point. The model was also calibrated to best available data from that time when statins did not have as much information on long-term effectiveness. Given that postmarket knowledge of statin effectiveness is greater now than in 2009, these results express uncertainty where greater knowledge now exists.

Report Results

The MDP model also unifies results across the three risk models, where there is noticeable variability in recommended treatment between studies. The Framingham model determines never to initiate statins for three of the nine metabolic states. The Archimedes risk model does not offer statin start points for all metabolic states, and predicts statin start points based on statistical inference rather than by generalizable samples of patients, making the model prone to statistical error.

In contrast, the UKPDS and Framingham risk models fit smoothed Weibull distributions across a well-defined population sample. The UKPDS and the Framingham model give numerically different, but qualitatively similar, optimal statin start time results. However, using the Archimedes risk model in the optimization did not produce a smooth pattern for initiating statin therapy as observed with the UKPDS and Framingham models. The authors attribute this to "statistical error" associated with the Archimedes estimates.

The study demonstrates the value of the MDP framework, providing insight into when to start statin treatment. As one would expect, the model generally shows that statins should be started earlier for more severely ill patients. Exactly how early

depends on the severity of the patient's condition, but also on model parameters and which risk model is used. Interestingly, for less severe and elderly patients, from the results of Figure in the article [3], it seems that it may not be worthwhile starting statin therapy at all. Women are in general recommended to start statin treatment later than men.

Decision Making

The study is an example of how the MDP modeling approach can provide personalized and clinically relevant recommendations (for patients of type x, start statins at age y) and integrate and compare different data sources and risk models. As there are many questions about the right time to start, stop, and switch treatment in medical care, this seems an underused and highly promising framework for economic evaluation.

Due to the dynamic nature of the MDP modeling, it provides the ability to personalize decisions, as opposed to one-size-fits-all policies and guidelines established for medical decisions. However, similar to other approaches, MDPs have assumptions based on data and/or the structure of the model. Once the results are obtained, sensitivity analyses can be performed (e.g., for some range of variation in the transition probabilities). Once satisfied with the solution, translation is in the form of guidelines and/or decision tools. Owing to the modeling and computational nature of the MDPs, they can easily be translated into decision support systems to use in practice.

This example showed the use of MDP for optimizing the start time of statin therapy. MDPs can be used for other similar decision-making problems for breast or prostate cancer screening, the decision for biopsy, initiating HIV therapy treatment policies, etc. The underlying theme is focusing on decisions over time, with decisions at one point affecting future states and decisions operating under constrained resources. The results of the optimization models can help establish optimal clinical guidelines [36].

Conclusion

In this second report, the task force's primary objective is to provide an overview of areas where optimization methods can be applied and describe three case studies illustrating the application of constrained optimization methods to critical clinical and health policy questions. The cases illustrate several major variants of these methods and demonstrate their potential in complementing the classical economic evaluation, decision-making framework.

In the first case study, linear programming methods were used to identify the optimal mix of HPV vaccination and screening to minimize the number of cervical cancer cases subject to a budget constraint. Similarly, in the second case study, MDP and dynamic programming were used to identify the optimal time to initiate statin therapy in type 2 diabetes patients. The first two case studies describe the translation of the original problem into its mathematical formulation as well as its estimation, interpretation, and use. In contrast, the third is an educational case that allows the reader to work through the formulation of a constrained optimization problem using the ISPOR Constrained Optimization Good Practice Checklist.

The health care sector faces major challenges with regard to appropriate diagnosis and treatment, allocation of scarce resources, design of policies, etc. These methods provide an approach for finding optimal solutions to complex problems in the face of constraints. As such, they are complementary to and build on the health economic models and simulation methods that are widely used to guide clinical and policy decision making.

Constrained optimization methods can improve the current reimbursement decision-making processes, which take budget constraints partially into account. In the constrained optimization framework, budget constraints can be incorporated explicitly, together with other types of constraints, like human resource or geographical equity constraints. In addition, when there are numerous treatment options available for treating patients with a specific condition, constrained optimization might prove to be an efficient method for developing treatment protocols or guidelines compared to the classical economic evaluation framework.

In the current health care landscape, health economic modeling is widely used to make reimbursement decisions for new technologies, particularly outside the United States.

Constrained optimization methods can help decision makers incorporate related considerations beyond the reimbursement decision itself, including the best way to integrate the new technology with the health care delivery system, as well as in technology disinvestment decisions. These are becoming crucial as personalized medicine and performance-based payment concepts become more common.

It is important to recognize that application of constrained optimization methods in health care is still an emerging area and that there are some challenges that must be addressed. Constrained optimization methods can be limited by data availability and quality, and validating an optimization model can be challenging. Choosing and applying the appropriate method can be difficult and require specific expertise. Interpreting results and knowing which solution algorithm is likely to be best require a level of methodological understanding and sophistication.

However, despite these obstacles, the application of constrained optimization methods to health care decision making offers substantial potential benefits, which make them a valuable addition to the arsenal of analytic methods at the disposal of the researcher. Approaching a problem in the context of mathematical optimization forces modelers to identify and quantify the endpoint that they are trying to accomplish. Most importantly, constrained optimization takes into account the limits placed on the solution by real-world factors, such as budgets, availability of treatments, staffing capacity, and patient characteristics. As a result, implementation of the identified optimal solution is much more likely to be feasible.

In a disease management problem, by treating patients optimally we have the potential to improve population health and enhance the value associated with health care expenditure. For individual patients, this means providing treatment with the proper therapy faster. For physicians, this can help provide optimal health outcomes for their patients, enhance the performance of their medical practice, and offer more efficient health care delivery. The task force hopes that these two reports will encourage modelers to explore the use of optimization methods and looks forward to seeing more published optimization applications and the development of further guidelines and resources as the use of these methods becomes more widespread.

Acknowledgments

The coauthors thank all those who commented orally during four task force workshop and forum presentations at ISPOR conferences in the United States and Europe in 2016 and 2017. The coauthors gratefully acknowledge the following 26 reviewers who contributed their time and expertise through submission of written comments on the ISPOR Optimization Methods Emerging Good Practices Task Force Reports 1 and 2: Ekkehard Beck, Bjorn Berg, J. Jaime Caro, Koen Degeling, Brian Denton, Beth Devine, Stephanie Earnshaw, Talitha Feenstra, Chris Hane, Katherine Hicks, Julie L. Higle, Ng Chin Hui, Adam Letchford, Dawn Lee,

Nan Liu, Lena Burgos Liz, Nikolas Popper, Farhan Abdul Rauf, Emily S. Reese, Ajit Singh, Eunju Todd, Pepijn Vemer, Amir Viyanchi, Richard J. Willke, Beth Woods, Greg Zaric. Their generous feedback improved the manuscript and made it an expert consensus ISPOR Task Force Report.

Special thanks to Bhash Naidoo, Senior Technical Adviser (Health Economics)–NICE Centre for Guidelines, National Institute for Health and Care Excellence, London, England, UK, for his comments.

We are especially grateful to ISPOR staff, including Elizabeth Molsen-David, for helping us get these task force reports done–from beginning to end–and to Kelly Lenahan for her assistance in producing these reports. We also thank Petrine Cerri for her help with reviewing, formatting, and compiling the references for this report.

Financial support: None of the authors received financial support for their participation in this ISPOR Task Force. All authors volunteered their time for the discussion, research, and writing of this report. This research was supported in part by ISPOR, which contributed one staff liaison for this project.

Supplemental Materials

Supplementary data associated with this article can be found in the online version at https://doi.org/10.1016/j.jval.2018.05.003.

REFERENCES

[1] Crown W, Buyukkaramikli N, Thokala P, et al. Constrained Optimization Methods in Health Services Research–An Introduction: Report 1 of the ISPOR Optimization Methods Emerging Good Practices Task Force. Value Health 2017;20:310–9.
[2] Demarteau N, Breuer T, Standaert B. Selecting a mix of prevention strategies against cervical cancer for maximum efficiency with an optimization program. Pharmacoeconomics 2012;30:337–53.
[3] Denton BT, Kurt M, Shah ND, et al. Optimizing the Start Time of Statin Therapy for Patients with Diabetes. Med Decis Making 2009;29:351–67.
[4] Stinnett AA, Paltiel AD. Mathematical programming for the efficient allocation of health care resources. Health Econ 1996;15:641–53.
[5] Airoldi M, Morton A, Smith JA, Bevan G. STAR–people-powered prioritization: a 21st-century solution to allocation headaches. Med Decis Making 2014;34:965–75.
[6] Martello S, Toth P. Knapsack problems: algorithms and computer implementations. New York, NY: John Wiley & Sons, Inc, 1990.
[7] Weinstein M, Zeckhauser R. Critical ratios and efficient allocation. Journal of Public Economics 1973;2:147–57.
[8] Thokala P, Dixon S, Jahn B. Resource modeling: the missing piece of the HTA jigsaw? Pharmacoeconomics 2015;33:193–203.
[9] Castillo-Chavez C, Feng Z. Global stability of an age-structure model for TB and its applications to optimal vaccination strategies. Math Biosci 1998;151:135–54.
[10] Juusola JL, Brandeau ML. HIV Treatment and Prevention: A Simple Model to Determine Optimal Investment. Med Decis Making 2016;36:391–409.
[11] Lee EK, Yuan F, Pietz FH, et al. Vaccine Prioritization for Effective Pandemic Response. Interfaces 2015;45:425–43.
[12] Marshall DA, Burgos-Liz L, IJzerman MJ, et al. Selecting a dynamic simulation modeling method for health care delivery research–part 2: report of the ISPOR Dynamic Simulation Modeling Emerging Good Practices Task Force. Value Health 2015;18:147–60.
[13] Marshall DA, Burgos-Liz L, IJzerman MJ, et al. Applying dynamic simulation modeling methods in health care delivery research–the SIMULATE checklist: treport of the ISPOR simulation modeling emerging good practices task force. Value Health 2015;18:5–16.
[14] Dimitrov NB, Meyers LA. In: Hasenbein JJ,ed. Mathematical Approaches to Infectious Disease Prediction and Control, 7. INFORMS TutORials in Operations Research; 2010. p. 1–25: ISBN 978-0-9843378-0-4. http://dx.doi.org/10.1287/educ.1100.0075.
[15] Roth AE, Sotomayor M. Two-sided matching. In: Aumann RJ, Hart S,eds. Handbook of Game Theory with Economic Applications. Amsterdam, The Netherlands: Elsevier Science B. V, 1992.
[16] Bertsimas D, Farias VF, Trichakis N. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. Operations Research 2013;61:73–87.
[17] Puterman ML. Markov Decision Processes: Discrete Stochastic Dynamic Programming. New York, NY: John Wiley & Sons, Inc, 2014.
[18] Shechter SM, Bailey MD, Schaefer AJ, Roberts MS. The optimal time to initiate HIV therapy under ordered health states. Operations Research 2008;56:20–33.
[19] Chhatwal J, Alagoz O, Burnside ES. Optimal breast biopsy decision-making based on mammographic features and demographic factors. Operations Research 2010;58:1577–91.
[20] Tosh J, Stevenson M, Akehurst R, Strong M. Simulation optimisation of Treatment Sequences for Rheumatoid arthritis. Value Health 2015;18:A343.
[21] Lee EK, Wu TL. Disease Diagnosis: Optimization-Based Methods. In: Floudas CA, Pardalos PM,eds. Encyclopedia of Optimization. Boston, MA: Springer US, 2009.
[22] Liberatore MJ, Nydick RL. The analytic hierarchy process in medical and health care decision making: A literature review. Eur J Oper Res 2008;189:194–207.
[23] Ehrgott M, Holder A, Reese J. Beam selection in radiotherapy design. Linear Algebra Appl 2008;428:1272–312.
[24] Lee EK, Zaider M. Operations research advances cancer therapeutics. Interfaces 2008;38:5–25.
[25] Bertsimas D, O'Hair A, Relyea S, Silberholz J. An Analytics Approach to Designing Combination Chemotherapy Regimens for Cancer. Management Science 2016;62:1511–31.
[26] Thokala P, Ochalek J, Leech AA, Tong T. Cost-Effectiveness Thresholds: the Past, the Present and the Future. Pharmacoeconomics 2018;36:509–22.
[27] Lin RC, Sir MY, Pasupathy KS. Multi-objective simulation optimization using data envelopment analysis and genetic algorithm: Specific application to determining optimal resource levels in surgical services. Omega 2013;41:881–92.
[28] Fu MC. Optimization for Simulation: Theory vs. Practice. INFORMS J Comput 2002;14:192–215.
[29] Barton RR. Metamodeling: a state of the art review. Proceedings of the 1994 Winter Simulation Conference. Orlando, FL, 1994, pp. 237–244. https://dx.doi.org/10.1109/WSC.1994.717134.
[30] Sargent RG. Verification and validation of simulation models, Proceedings of the 2009 Winter Simulation Conference (WSC), Austin, TX, 2009, pp. 162–176. https://doi.org/10.1109/WSC.2009.5429327. Available from: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5429327&isnumber=5429163.
[31] Law AM. How to Build Valid and Credible Simulation Models. In: Perrone LF, Wieland FP, Liu J, Lawson BG, Nicol DM, Fujimoto RM,eds. Proceedings of the 2006 Winter Simulation Conference. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc., 2006. p. 58–66.
[32] Eddy DM, Hollingworth W, Caro JJ, et al. Model transparency and validation: a report of the ISPOR-SMDM Modeling Good Research Practices Task Force-7. Med Decis Making 2012;32:733–43.
[33] Vemer P, Corro Ramos I, van Voorn GA, et al. AdViSHE: A Validation-Assessment Tool of Health-Economic Models for Decision Makers and Model Users. Pharmacoeconomics 2016;34:349–61.
[34] Spall JC. Introduction to Stochastic Search and Optimization: Estimation. Simulation, and Control. New York, NY: John Wiley & Sons, Inc, 2003.
[35] Gorissen BL, Yanikoglu I, den Hertog D. A practical guide to robust optimization. Omega 2015;53:124–37.
[36] Steimle LND, Brian T. Markov Decision Processes for Screening and Treatment of Chronic Diseases. In: Boucherie RJ, van Dijk NM,eds. Markov Decision Processes in Practice. Basel, Switzerland: Springer International Publishing, 2017.