Socio-Economic Planning Sciences 47 (2013) 172–182

Contents lists available at SciVerse ScienceDirect

Socio-Economic Planning Sciences

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

Use of a Markov decision process model for treatment selection in an asymptomatic disease with consideration of risk sensitivity

Vera Tilson*, David A. Tilson

William G. Simon Graduate School of Business, University of Rochester, Rochester, NY 14627, USA

ARTICLE INFO

Article history:
Available online 27 September 2012

Keywords:
Medical decision making
Markov decision processes in healthcare

ABSTRACT

Some potentially dangerous diseases are completely asymptomatic. Their diagnosis as incidental findings of ever-more-sensitive medical imaging can leave patients and physicians in something of a quandary. The patient feels well, and potential interventions to stave off long-term deterioration or death bring with them immediate risks. We discuss the use of a Markov Decision Process (MDP) model (rather than Monte Carlo simulation of a Markov Model) to create a tool for analyzing individual treatment decisions for asymptomatic chronic diseases where a patient's condition cannot improve. We formulate a finite-horizon MDP model to determine optimal treatment plans and discuss three distinct optimality criteria: (a) maximizing expected quality-adjusted-life years with and without discounting, (b) maximizing the expected number of life years in good health, and (c) maximizing the expected utility for number of years in good health. In (c) we assume exponential utility and consider different risk aversion factors reported in the medical literature. We illustrate the model's use by considering asymptomatic intracranial aneurysm. Our model builds on a simulation model [19] created to examine treatment recommendations based on cost-effectiveness. We demonstrate that incorporating risk aversion leads to "no treatment" recommendations for some types of aneurysm. Furthermore, the use of alternate patient-selected criteria leads to recommendations that vary from [19] in several scenarios. We also discuss the use of the software as a decision support tool to help make individualized treatment recommendations and demonstrate that the computational performance of the algorithm makes its use feasible during a short office visit.

1. Introduction

The American Recovery and Reinvestment Act of 2009 authorized $1.1 billion in support of comparative effectiveness research in medicine. According to the U.S. Department of Health and Human Services website, "comparative effectiveness research provides information on the relative strengths and weakness of various medical interventions. Such research will give clinicians and patients valid information to make decisions that will improve the performance of the U.S. health care system."

Even though comparative effectiveness is not tantamount to cost-effectiveness, in many societal contexts treatment options are compared on the basis of expected cost per quality-adjusted life year (QALY). A common methodology for such comparisons combines Monte Carlo simulation with a Markov model of the disease. A search of PubMed database for articles with terms "Markov" and "cost-effectiveness" in the title or abstract returns close to 1600 papers added to the database since 1986. Markov model definition includes the specification of patient health states, the model cycle – time between state transitions – and the time horizon. *Transition probabilities* between various states due to treatments or disease progression are obtained from medical literature, as are *health state utilities*, and the *costs* associated with the disease or its treatment. After specifying the model, analysts use Monte Carlo simulation to obtain the expected discounted number of quality-adjusted life years as well as expected discounted costs for a simulated cohort of patients. Treatments are compared based on the ratio of these expected discounted costs and QALYs.

Treatment recommendations based on cost-effectiveness may not be valid across geographies since the magnitude of the financial costs, and the entities bearing them, vary so widely. Even comparing treatments based on QALY alone, without regard to costs, is problematic, since there is neither agreement on what populations should be surveyed to obtain the quality-of-life utilities nor on the methods and survey instruments that should be

* Corresponding author. Tel.: +1(585)275 7956.
E-mail addresses: vera.tilson@simon.rochester.edu (V. Tilson), david.tilson@simon.rochester.edu (D.A. Tilson).

0038-0121/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.seps.2012.09.003

used [8,15]. Similarly, the selection of discount factors used to model preferences in health over time is also debated [18]. Finally, even if the recommendations were based on *expected* QALYs arising from *individual* quality-of-life weights and the disease probabilities for individual patients, they would still fail to take patients' risk attitudes into consideration. As noted by Asch and Hershey [4], taking a patient's risk attitude into consideration is important because no single patient will have the opportunity to try different treatments multiple times, so it is unreasonable to conjecture that individual preferences would be based on expected outcomes.

Patients can differ in their manifestation of a disease, and the most effective treatment differs both by the particular manifestation and an individual patient's attitudes to various dimensions of treatment outcomes as well as risk preferences. Online decision support tools (e.g., http://www.lifemath.net/cancer/) help patients analyze probabilistic benefits of different therapies. In this paper, we discuss the use of a Markov Decision Process (MDP) model (in contrast to Monte Carlo simulation of a Markov Model) to create a tool for analyzing individual treatment decisions with consideration of a patient's risk attitude. To date, MDP analysis has seen little use in medical decision making [2], possibly because the analysis requires the development of custom software, while off-the-shelf software packages are available to perform Monte Carlo simulation with Markov Models in medical settings (e.g., TreeAge). Many treatment decisions are sequential decision problems, and MDP models are powerful tools for analyzing such problems. Unlike simulation models, MDPs are *efficient* at identifying *optimal* treatment policies under a variety of reasonable objectives, such as, (i) maximizing expected QALYs with quality-of-life scores obtained from the individual patient, (ii) maximizing the expected number of years in good health, (iii) maximizing the expected utility of the number of years in good health, and (iv) maximizing probability of being well at a particular time in the future, etc. The two latter criteria encompass risk sensitivity, acknowledged as an important consideration in medical decisions with complex tradeoffs (e.g. [4,5]). Even when a treatment decision problem is relatively simple and consists only of identifying the optimal treatment time and type, algorithms for solving MDPs are computationally efficient relative to simulation models and therefore more suitable for utilization in decision aids that can be used during a patient's visit with a physician.

We focus our analysis on a particular type of disease: asymptomatic chronic diseases diagnosed as an incidental finding. The disease may or may not progress, but it can be treated with a risky intervention prior to any deterioration in the patient's condition. Advances in medical imaging and its more frequent use have resulted in increased asymptomatic incidental findings. One such condition, asymptomatic intracranial aneurysm, is an abnormal bulging outward of one of the brain's arteries. An aneurysm's rupture can lead to stroke, brain damage or death. However, an aneurysm may never rupture, and an asymptomatic aneurysm creates no physical discomfort for a patient. The dilemma for the patient is that the available treatments are risky and provide no immediately apparent benefit, as, at the time of the diagnosis, the patient is otherwise well[1]. We specifically discuss incorporating a patient's risk attitude into the MDP objective.

Takao and Noji [19] proposed a Markov model of intracranial aneurysm in their cost-effectiveness study. The model specifies upwards of twenty probability parameters related to the disease and its treatment. It is doubtful that any physician or patient can determine the joint implication of so many parameters without the benefit of a decision support system. Additionally, Walling [20] wrote that the optimal management of unruptured intracranial aneurysms is "highly controversial because of uncertainty about the probability of rupture and the risks of surgical repair." Despite the significant uncertainty surrounding the disease and treatment-related probabilities, physicians make treatment recommendations based on a variety of criteria and rules of thumb. Another advantage of MDP analysis is that a suitable software implementation allows for rapid sensitivity analysis, to determine whether uncertainty about parameter values would affect treatment choice.

This paper is organized as follows: In Section 2, we formulate a general MDP model for a disease and discuss risk-sensitive objective formulation for an asymptomatic disease. In the following sections, we use the example of intracranial aneurysm to illustrate the approach. Model specification for asymptomatic intracranial aneurysms and their treatment is presented in Section 3. Section 3.3 presents sample results of a quantitative analysis for intracranial aneurysm, showing how treatment recommendations differ by optimization criteria. Section 4 concludes with a discussion of the described approach's contribution to the medical decision making literature.

2. Model formulation

We define a patient's state s as a duple: $\{a,\sigma\}$. Element a is the patient's age, and $\sigma \in \Sigma$, where Σ is the set of health states. A health state combines the patient's physical well being (e.g., well, disabled) and eligibility for treatment as well as his possible deterioration due to the disease. Let $\hat{\sigma} \in \Sigma$ designate the initial state at the time of diagnosis, i.e., the state where the patient is well and has not yet received any treatment. Let $Y \in \Sigma$, be the set of all states where the patient is well, i.e., asymptomatic.

The decision epochs in the finite-horizon model are numbered with integers 0 through N. At the start of epoch 0, the patient's current age is denoted by a_0. We assume that from epoch n to epoch $n+1$ the patient ages by c, where c is the model's *cycle length*. We further assume that patients do not live past the age of A. So, N is the smallest integer larger than or equal to $(A-a_0)/c$, that is $N = \lceil (A-a_0)/c \rceil$.

Patients transition from a state $s_n = \{a_n, \sigma_n\}$ in epoch n, to state $s_{n+1} = \{a_{n+1}, \sigma_{n+1}\}$ in epoch $n+1$. While the transition of the age element is deterministic:

$$a_{n+1} = a_n + c, \tag{1}$$

the transition from σ_n to σ_{n+1} is probabilistic.

The purpose of any treatment is to prevent an irreversible deterioration of the patient's health. The treatment can only be rendered when the patient is well. Further, we assume that at most one interventional treatment is possible over a lifetime, so the patient is deciding between different treatments and the time to administer them. Thus T_s, the set of feasible treatments depending on the patient's state s is defined as

$$T_s = \begin{cases} \{\text{no treatment}, \text{intervention } 1, \ldots, \text{intervention } k\} \text{ if } \sigma = \hat{\sigma} \\ \{\text{no treatment}\} \text{otherwise} \end{cases} \tag{2}$$

A treatment *policy* $\boldsymbol{t} = \{t_0, t_1, \ldots, t_{N-1}\}$ prescribes a treatment t_n for a patient of age a_n in health state $\sigma = \hat{\sigma}$.

Let π_σ^n denote the probability of a patient being in health state σ at age a_n. The vector $\boldsymbol{\pi}^n = \{\pi_0^n, \pi_1^n, \ldots, \pi_{\hat{\sigma}}^n\}$[2] is a function of a selected treatment policy found as

[1] Note that we refer to an asymptomatic patient with an intracranial aneurysm as "well" following the convention we observed in the medical literature.

[2] Without loss of generality, we assume that $\hat{\sigma}$ is the highest numbered state in $\sum$.