

# Chapter 24

## Markov Models and Cost Effectiveness Analysis: Applications in Medical Research

Matthieu Komorowski and Jesse Raffa

### Learning Objectives

Understand how Markov models can be used to analyze medical decisions and perform cost-effectiveness analysis.

This case study introduces concepts that should improve understanding of the following:

1. Markov models and their use in medical research.
2. Basics of health economics.
3. Replicating the results of a large prospective randomized controlled trial using a Markov Chain and Monte Carlo simulations, and
4. Relating quality-adjusted life years (QALYs) and cost of interventions to each state of a Markov Chain, in order to conduct a simple cost-effectiveness analysis.

### 24.1 Introduction

Markov models were initially theroreticized at the beginning of the 20th century by Russian mathematician Andrey Markov [1]. They are stochastic processes that undergo transitions from one state to another. Over the years, they have found countless applications, especially for modeling processes and informing decision making, in the fields of physics, queuing theory, finance, social sciences, statistics and of course medicine. Markov models are useful to model environments and **problems involving sequential, stochastic decisions over time**. Representing such environments with decision trees would be confusing or intractable, if at all possible, and would require major simplifying assumptions [2]. Markov models can be examined by an array of tools including linear algebra (brute force), cohort simulations, Monte Carlo simulations and, for Markov Decision Processes, dynamic programming and reinforcement learning [3, 4].

A fundamental property of all Markov models is their **memorylessness**. They satisfy a first-order **Markov property** if the probability to move a new state to  $s_{t+1}$  only depends on the current state  $s_t$ , and not on any previous state, where  $t$  is the current time. Said otherwise, given the present state, the future and past states are independent. Formally, a stochastic process has the first order Markov property if the conditional probability distribution of future states of the process (conditional on both past and present values) depends only upon the present state:

$$P(s_{t+1}|s_1, s_2, \dots, s_t) = P(s_{t+1}|s_t)$$

This chapter will provide a brief introduction to the most common Markov models, and outline some potential applications in medical research and health economics. The last section will discuss a practical example inspired from the medical literature, in which a Markov chain will be used to conduct the cost-effectiveness analysis of a particular medical intervention. In general, the crude results of a study are unable to provide the necessary information to fully implement cost-effectiveness analysis, thus demonstrating the value of expressing the problem as a Markov Chain.

## 24.2 Formalization of Common Markov Models

The four most common Markov models are shown in Table 24.1. They can be classified into two categories depending or not whether the entire sequential state is observable [5]. Additionally, in Markov Decision Processes, the transitions between states are under the command of a control system called the agent, which selects actions that may lead to a particular subsequent state. By contrast, in Markov chains and hidden Markov models, the transition between states is autonomous. All Markov models can be finite (discrete) or continuous, depending on the definition of their state space.

### 24.2.1 The Markov Chain

The discrete time Markov chain, defined by the tuple  $\{S, T\}$  is the simplest Markov model, where  $S$  is a finite set of states and  $T$  is a state transition probability matrix,

**Table 24.1** Classification of Markov models

	Fully observable system	Partially observable systems
Autonomous system	Markov chain (MC)	Hidden Markov model (HMM)
System containing a control process	Markov decision process (MDP)	Partially observable Markov decision process (POMDP)



**Fig. 24.1** Example of a Markov chain, defined by a set  $S$  of finite states {Healthy, Ill} and a transition matrix, containing the probabilities to move from current state  $s$  to next state  $s'$  at each iteration

**Table 24.2** Example of a transition matrix corresponding to Fig. 24.1

		Next state $s$		Total
		Healthy	Ill	
Initial state $s$	Healthy	0.9	0.1	1
	Ill	0.5	0.5	1

$T(s', s) = P(s_{t+1} = s' | s_t = s)$ . A Markov chain can be **ergodic**, if it is possible to go from any state to every other state in finitely many moves. Figure 24.1 shows a simple example of a Markov Chain.

In the transition matrix, the entries in each column are between 0 and 1 (inclusive) and their sum is 1. Such vectors are called **probability vectors**. The Table 24.2 shows the transition matrix corresponding to Fig. 24.1. A state is said to be **absorbing** if it is impossible to leave it (e.g. death).

### 24.2.2 Exploring Markov Chains with Monte Carlo Simulations

Monte Carlo (MC) simulations are a useful technique to explore and understand phenomena and systems modeled under a Markov model. MC simulation generates pseudorandom variables on a computer in order to approximate difficult to estimate quantities. It has wide use in numerous fields and applications [6]. Our focus is on the MC simulation of a Markov chain, and it is straightforward once a transition probability matrix,  $T(s', s)$ , and final time  $t^*$  have been defined. We will assume at the index time ( $t = 0$ ), the state is known, and call it  $s_0$ . At  $t = 1$ , we simulate a categorical random variable using the  $s_0$ th row of the transition probability matrix  $T(s', s)$ . We repeat this  $t = 1, 2, \dots, t^* - 1, t^*$  to simulate *one simulated instance* of the Markov chain we are studying. One simulated instance only tells us about one possible sequence of transitions out of very many for this Markov chain, and we need to repeat this many ( $N$ ) times, recording the sequence of states for each of the simulated instances. Repeating this process many times, allows us to estimate quantities such as: the probability at  $t = 5$ , that the chain is in state 1; the average

proportion of time spent in state 1 over the first 10 time points; or the average length of the longest consecutive streak in state 1 in the first  $t^*$  time points.


Using the example shown in Fig. 24.1, we will estimate the probability for someone to be healthy or ill in 5 days, knowing that he is healthy today. MC methods will simulate a large number of samples (say 10,000), starting in  $s_0 = \text{Healthy}$  and following the transition matrix  $T(s', s)$  for 5 steps, sequentially picking transitions to  $s'$  according to their probability. The output variable (the value of the final state) is recorded for each sample, and we conclude by analyzing the characteristics of the distribution of this output variable (Table 24.3).

The distribution of the final state at day + 5 for 10,000 simulated instances is represented on Fig. 24.2.

Table 24.4 reports some sample characteristics for “healthy” state on day 5 for 100 and 10,000 simulated instances, which illustrates why it is important to simulate a very large number of samples.

**Table 24.3** Example of health forecasting using Monte Carlo simulation

	Instance 1	Instance 2	...	Instance 10,000
Today	Healthy	Healthy	...	Healthy
Day + 1	Healthy	Healthy		Healthy
Day + 2	Healthy	Ill		Healthy
Day + 3	Healthy	Ill		Ill
Day + 4	Healthy	Ill		Healthy
Day + 5	Healthy	Ill	...	Healthy



**Fig. 24.2** Distribution of the health on day 5, for 10,000 instances

**Table 24.4** Sample characteristics for 100 and 10,000 simulated instances


	100 simulated instances	10,000 simulated instances
Mean	0.81	0.83
Standard deviation	0.39	0.37
95 % confidence interval for the mean	0.73–0.89	0.83–0.84

By increasing the number of simulated instances, we drastically increase our confidence that the true sample mean falls within a very narrow window (0.83–0.84 in this example). The true mean calculated analytically is 0.838, which is very close to the estimate generated from MC simulation.

### 24.2.3 Markov Decision Process and Hidden Markov Models

Markov Decision Processes (MDPs) provide a framework for running reinforcement learning methods. MDPs are an extension of Markov chains, which include a control process. MDPs are a powerful and appropriate technique for modeling medical decision [3]. MDPs are most useful in classes of problems involving **complex, stochastic and dynamic decisions like medical treatment decisions**, for which they can find optimal solutions [3]. Physicians will always need to make subjective judgments about treatment strategies, but mathematical decision models can provide insight into the nature of optimal choices and guide treatment decisions.

In Hidden Markov models (HMMs), the state space is only partially observable [7]. It is formed by two dependent stochastic processes (Fig. 24.3). The first is a classical Markov chain, whose states are not directly observable externally, therefore “hidden.” The second stochastic process generates observable emissions, conditional on the hidden process. Methodology has been developed to decode the hidden states from the observed data and has applications in a multitude of areas [7].



**Fig. 24.3** Example of a hidden Markov model (HMM)

### 24.2.4 *Medical Applications of Markov Models*

MDPs have been praised by authors as being a powerful and appropriate approach for modeling sequences of medical decisions [3]. Controlled Markov models can be solved by algorithms such as dynamic programming or reinforcement learning, which intends to identify or approximate the optimal policy (set of rules that maximizes the expected sum of discounted rewards).

In the medical literature, Markov models have explored very diverse problems such as timing of liver transplant [8], HIV therapy [9], breast cancer [10], Hepatitis C [11], statin therapy [12] or hospital discharge management [5, 13]. Markov models can be used to describe various health states in a population of interest, and to detect the effects of various policies or therapeutic choices. For example, Scott et al. has used a HMM to classify patients into 7 health states corresponding to side effects of 2 psychotropic drugs [14]. The transitions were analyzed to specify which drug was associated with the least side-effects. Very recently, a Markov chain model was proposed to model the progression of diabetic retinopathy, using 5 pre-defined states, from mild retinopathy to blindness [15]. MDPs have also been exploited in medical imaging applications. Alterovitz has used very large MDPs (800,000 states) for motion planning in image-guided needle steering [16].

Besides those medical applications, Markov models are extensively used in health economics research, which is the focus of the next section of this chapter.

## 24.3 Basics of Health Economics

### 24.3.1 *The Goal of Health Economics: Maximizing Cost-Effectiveness*

This section provides the reader with a minimal background about health economics, followed by a worked example. Health economics intends to maximize “value for money” in healthcare, by optimizing not only clinical effectiveness, but also cost-effectiveness of medical interventions. As explained by Morris: “Achieving ‘value for money’ implies either a desire to achieve a predetermined objective at least cost or a desire to maximise [sic] the benefit to the population of patients served from a limited amount of resources” [17].

Two main approaches can be outlined in health economics: cost-minimization and cost-effectiveness analysis (CEA). In both cases, the purpose is identical: to identify which treatment option is the most cost-effective. Cost minimization deals with the simple case where the several treatment options available have the same effectiveness but different costs. Quite logically, cost-minimization will favor the cheapest option. CEA represents a more likely scenario and is more widely used.

In CEA, several options with different costs and different effectiveness are compared. The analysis will compute the relative cost of an improvement in health, and metrics to optimally inform decision makers.

### 24.3.2 Definitions

#### ***Measuring Outcome: Survival, Quality of Life (QoL), Quality-Adjusted Life-Years (QALY)***

Outcomes are assessed in terms of enhanced survival (“*adding years to life*”) and enhanced quality of life (QoL) (“*adding life to years*”) [17]. Although sometimes criticized, the concept of Quality-adjusted life-years (QALY) remains of central importance in cost-utility analysis [18]. QALYs apply weights that reflect the QoL being experienced by the patient. One QALY equates to one year in perfect health. Perfect health is equivalent to 1 while death is equivalent to 0. QALYs are estimated by various methods including scales and questionnaires filled by patients or external examiners [19]. As an example, the EuroQoL EQ 5D questionnaire assesses health in 5 dimensions: mobility, self-care, usual activities, pain/discomfort and anxiety/depression.

#### ***Cost-Effectiveness Ratio (CER)***

The cost-effectiveness ratio (CER) will inform the decision makers about the cost of an intervention, relative to the health benefits this intervention generates. For example, an intervention costing \$20,000 per patient and providing 5 QALYs (5 years of perfect health) has a CER of  $\$20,000/5 = \$4000$  per QALY. This measure allows a direct comparison of cost-effectiveness between interventions.

#### ***Incremental Cost-Effectiveness Ratio (ICER)***

The incremental cost-effectiveness ratio (ICER) is a measure very commonly reported in the health economics literature and allows comparing two different interventions in terms of “cost of gained effectiveness.” It is computed by dividing the difference in cost of 2 interventions by the difference of their effectiveness [20].

As an example, if treatment A costs \$5000 per patient and provides 2 QALYs, and treatment B costs \$8000 while providing 3 QALYs, the ICER of treatment B will be:

$$\frac{(\$8000 - \$5000)}{3 - 2} = \$3000$$

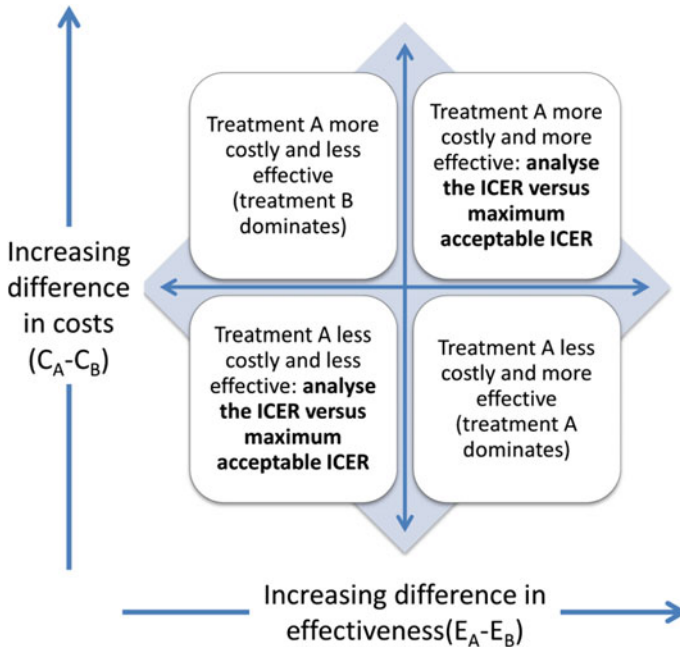
Said otherwise, it will cost \$3000 more to gain one more QALY with treatment B, for this particular medical condition. ICER can inform decision makers about the

need to adopt or fund a new medical intervention. Schematically, if the ICER of a new medical intervention lies below a certain threshold, it means that health benefits can be achieved with an acceptable level of spending.

### *The Cost Effectiveness Plane*

The cost-effectiveness plane (CE plane) is an important tool used in CEA (Fig. 24.4). It aims to clearly illustrate differences in costs and effects between different strategies, whether they comprise medical interventions, treatments, or even a combination of the two.

The CE plane consists of a four-quadrant diagram where the X-axis represents the incremental level of effectiveness of an outcome and the Y-axis represents the additional total cost of implementing this outcome. For example, the further right you move on the X-axis, the more effective the outcome. In the upper-right quadrant, a treatment may receive funding if its ICER lies below the maximum acceptable ICER threshold.



**Fig. 24.4** The cost-effectiveness plane, comparing treatment A with treatment B



## 24.4 Case Study: Monte Carlo Simulations of a Markov Chain for Daily Sedation Holds in Intensive Care, with Cost-Effectiveness Analysis

This example is inspired by the publication by Girard et al. [21], and will allow us to illustrate how to construct and examine a simple Markov Chain to represent a medical intervention, how to relate QALYs and cost of interventions to each state of the Markov Chain, in order to carry out a cost-effectiveness analysis. In this prospective randomized controlled trial, the authors evaluated the impact of daily sedation holds in intensive care on various outcomes such as the number of ventilator-free days, delirium and 28-day mortality. In the ICU, patients frequently undergo mechanical ventilation in the setting of severely impaired consciousness, after heavy surgical procedures, and when suffering from severe respiratory failure. Therapeutically, patients are sedated to maximize their comfort. A growing body of literature, however, has identified the risks of continuous sedation in the ICU, as it is associated with increased mortality, delirium, duration of mechanical ventilation and length of ICU and hospital stay [22]. To strike the right balance between maintaining sedation and mechanical ventilator support as long as the patient needs it, but also moving to extubation as soon as possible, Girard and colleagues proposed actively waking up the patients daily to assess their readiness to come off of the ventilator. The main results are shown in Table 24.5.


In this case study example, we will attempt to approximate those results using a very simple 3-state Markov Chain examined by MC simulation. As an exercise, we will extend the study to CEA. This tutorial will provide the reader with all the tools necessary to implement in other contexts Markov Chain MC simulation methods and simple cost-effectiveness studies.

Most of the study results can be approximated using a very crude 3-state Markov chain (Fig. 24.5), with the following state space: {Intubated, Extubated, Dead}. In this simplistic model, only 7 transitions are possible, and the state ‘dead’ is absorbing.

**Table 24.5** Main results from the original study

	Intervention group	Control group
Ventilator-free days (mean)	14.7	11.6
Ventilator-free days (median)	20.0	8.1
Patients Successfully extubated at 28 days (%)	≈93	≈88
28 day mortality (%)	29	35

**Fig. 24.5** The 3-state Markov chain used in this example



Two different transition matrices can be built by trial-and-error, corresponding to the intervention and control arms of the study (Table 24.6). They correspond to the daily probabilities of transitioning from one state to another. The initial values were selected using a few simple assumptions: the state ‘death’ is absorbing, the probability to remain intubated or extubated is larger than the probability to change state, the risk of dying while intubated is larger than when extubated, and the total of each row in the transition matrix is one. Another assumption is that the intervention (daily sedation hold) will change the probability of successful extubation and mortality, hence the transition matrix. After each modification, the number of patients in each state was computed for 28 days (results in Table 24.8), so as to try to match the initial study’s results as closely as possible.

We can check to see if our code is running correctly by comparing important aspects of the simulation to known theoretical properties of probability theory and Markov Chains. For example, in our example all patients are assumed to be intubated at  $t = 0$ . Under our Markov model, the waiting time until extubation or death can be determined theoretically, but how to determine this is beyond the scope of this chapter. This waiting time,  $W^*$ , is a discrete random variable with a geometric distribution. Geometric distributions have probability mass functions, for a given waiting time,  $w$  of  $p(w) = (1 - p)p^{(w-1)}$ , where  $p$  is the probability of remaining intubated. In Fig. 24.6, we compare the number of times we observed different values of  $w$  to what we would expect under the true theoretical distribution of  $W^*$ , by computing  $Np(w)$ , where  $N$  is the number of simulated instances we computed.

**Table 24.6** Transition matrices used in the case study

Intervention group		Next state S'		
		I	E	D
Initial state S	I	0.862	0.12	0.018
	E	0.0088	0.982	0.0092
	D	0	0	1
Control group		Next state S'		
		I	E	D
Initial state S	I	0.878	0.1	0.022
	E	0.01	0.978	0.012
	D	0	0	1


We can see that our simulation follows very closely to what is theoretically known to be true.

In order to perform CEA, each state must be assigned a value for QALYs and cost. For the purpose of this example, let's also assume the values for QALYs and daily costs shown in Table 24.7.

Table 24.8 shows the results of the first iterations for the control group, when starting with 100 patients intubated (*function* IED\_transition.m). At each time step, the number of patients still intubated corresponds to the patients who stayed intubated, minus the patients who became extubated (daily probability of 10 %) and those who died (probability of 2.2 %), plus the extubated patients who had to be re-intubated (probability 1 %). After 28 days, the cumulated mortality reaches 35.6 %, and the ratio of patients extubated among the patients still alive is 88.8 %, hence matching quite closely the results of the initial study. At each time step, the sum of the QALYs and costs for all the patients is computed, as well as their cumulative values. The number of QALYs initially increases as more patients become extubated, then decreases as a consequence the number of patients dying.

**Table 24.7** Definition of QALY and daily cost for each state

State	I	E	D
QALY	0.5	1	0
Daily cost (\$)	2000	1000	0



**Fig. 24.6** Example of the life expectancy in state “I” in the control group, with fitted geometric distribution. The bar chart represents the distribution of the time spent in the state “intubated” of the Markov chain, before transitioning to another state, for 5000 samples

**Table 24.8** Number of patients in each state, QALYs and cost analysis, during 28 iterations (control group)

Day	I	E	D	Extubated/Alive	QALYs	Cumulative QALYs	Daily cost (K \$)	Cumulative cost (K\$)
0	100.00	0.00	0.00	0.00	50.00	50.00	200.00	200
1	87.80	10.00	2.20	0.10	53.90	103.90	185.60	386
2	77.19	18.56	4.25	0.19	57.15	161.05	172.94	559
3	67.96	25.87	6.17	0.28	59.85	220.90	161.78	720
4	59.92	32.10	7.98	0.35	62.06	282.96	151.95	872
5	52.94	37.38	9.68	0.41	63.85	346.81	143.25	1016
...	...	...	...	...	...	...	...	...
28	7.19	57.21	35.60	0.89	60.80	1863.84	71.59	3184

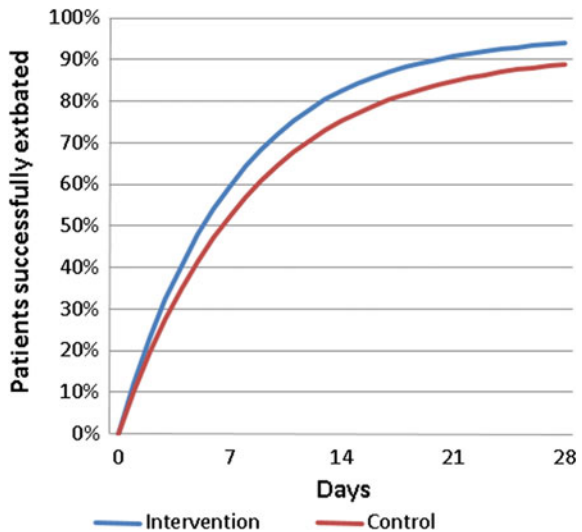
The following figure represents the ratio of number of patients extubated over number of patients alive, over time and for both strategies (Fig. 24.7). It can be compared to the original figure in the source article.

By simulating the distribution of the average number of ventilator-free days, and its characteristics, can be computed for both strategies (*function* `MCMC_solver.m`). The following Table 24.9 shows examples of patients' states computed using the transition matrix of the control group.


The distribution of ventilator-free days in our 10,000 samples is plotted shown in Fig. 24.8.

The mean and median number of ventilator-free days for both groups is shown in Table 24.10.

**Fig. 24.7** Modelled primary outcome of the study using a Markov chain



**Fig. 24.8** Ventilator-free days for 10,000 samples, for the intervention and control group



**Table 24.9** Computing the number of ventilator-free days by Monte Carlo (10,000 simulated instances)

Day	Instance 1	Instance 2	Instance 3	...	Instance 10,000
0	I	I	I		I
1	I	I	I		I
2	I	I	I		I
3	I	I	I		I
4	I	I	I		I
5	I	I	I		I
6	I	I	I		I
7	I	I	I		E
8	E	E	I		E
9	E	E	I		E
10	I	E	I		E
...	...	...	...		...
28	D	D	D		E
Total ventilator-free days	7	3	0	...	22

The cost-effectiveness ratio at 28 day of the both strategies can be computed by dividing the final cumulative cost by the cumulative QALYs (Table 24.11).

The intervention is more expensive but is also associated with health benefits (significantly more QALYs). It belongs to the upper-right quadrant of the CE plane,

**Table 24.10** Mean and median number of ventilator-free days for both groups

Number of ventilator-free days	Intervention group	Control group
Mean	17.1	15.9
Median	20	18

**Table 24.11** Cost-effectiveness ratio in both groups

	Intervention group	Control group
Cumulative cost (K\$)	3213	3184
Cumulative QALYs	2029	1864
Cost-effectiveness ratio (\$ per QALY)	1583	1708

where the ICER is used to determine the cost-effectiveness of an intervention. The ICER of this intervention is shown below:

$$ICER = \frac{(3,213,000 - 3,184,000)}{(2029 - 1864)} = 177.3$$

According to this crude analysis, Sedation holds appear to be a very cost-effective strategy, costing only \$177 more per additional QALY, relative to the control strategy. Reducing the value (QALY) of the state E from 1 to 0.6 significantly increases the ICER to \$1918 per QALY gained, demonstrating the huge impact that the definition of our health states has on the results of the CEA. Likewise, increasing the daily cost of state E from \$1000 to \$1900 (now only slightly cheaper than state I) leads to a much more expensive ICER of \$2041 per QALY gained. Some medical interventions may or may not be funded depending on the assumptions of the model!

## 24.5 Model Validation and Sensitivity Analysis for Cost-Effectiveness Analysis

An important component to any CEA is to assess whether the model is appropriate for the phenomena being examined, which is the purpose of model validation and sensitivity analyses. In the previous section, we model daily sedation hold as a Markov chain with a known transition probability matrix and costs. Deviations from this model can come in at least two types.

First, the use of a Markov Chain may be inappropriate to describe how subjects transition from the intubation, extubation and death states. It was presumed that this process follows a first-order Markov chain. Given enough real clinical data we can test to see if this assumption is reasonable. For example, given the transition probability matrices above, we can calculate quantities via MC simulation and

compare them to values reported in the real data. For instance, the authors report a 28-day mortality rate of 29 and 35 % in the intervention and control groups, respectively. From our simulation study, we estimate these quantities to be 27 and 35 %, which is reasonably close. One can perform formal goodness-of-fit testing as well to better assess if any differences noted provide any evidence that the model may be mis-specified. This process can also be repeated for other quantities, for example, the mean number of ventilator-free days.

In addition to validating the Markov model used to simulate the states and transitions for the system of interest, it is also important to perform a sensitivity analysis on the assumptions and parameters used in the simulation. Performing this step allows one to see how sensitive the results are to slight changes to parameter values. Choosing which parameters values to use in sensitivity analyses can be difficult, but some good practices are to find other parameters (e.g., transition probability matrices) reported in other studies of a similar type. For cost estimates, one may want to try costs reported in other countries, or incorporate important economic parameters like inflation. If using these other scenarios drastically affects the conclusions drawn from the simulation study, this does not necessarily mean that the study was a failure, but rather that there are limits to the generalizability of the simulation study's results. If particular parameters cause great fluctuations this may warrant further investigation into why this is the case. In addition to changing the parameters, one may try to alter the model significantly, by for example, using a higher order Markov model or semi-Markov model in place of a simple first order assumption, but these are advanced topic beyond the scope of this chapter.

The theoretical concepts introduced in the first sections of this chapter were applied to a concrete example coming from the medical literature. We demonstrated how clinical states and transition probabilities could be defined ad hoc, and how the stationary distribution of the chain could be estimated using Monte Carlo methods. The methodology outlined in this chapter will allow the reader to expand the results of other interventional studies to CEA, but countless other applications of Markov models exist, in particular in the domain of decision support systems.

## 24.6 Conclusion

Markov models have been used extensively in the medical literature, and offer an appealing framework for modeling medical decision making, with potential powerful applications in decision support systems and health economics analysis. They represent relatively simple mathematical models that are easy to grasp by non-data scientists or non-statisticians. Very careful attention must be paid to the verification of a fundamental assumption which is the Markov property, without which no further analysis should be carried out.

## 24.7 Next Steps

This tutorial hopefully provided basic tools to understand or develop CEA and Markov chains to model the effect of medical interventions. For more information on health economics, the reader is directed towards external references, such as the work by Morris and colleagues [17]. Guidance regarding the use of more advanced Markov models such as MDPs and HMMs is beyond the scope of this book, but numerous sources are available, such as the excellent Sutton and Barto, freely available online [4].

**Open Access** This chapter is distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, duplication, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, a link is provided to the Creative Commons license and any changes made are indicated.

The images or other third party material in this chapter are included in the work's Creative Commons license, unless indicated otherwise in the credit line; if such material is not included in the work's Creative Commons license and the respective action is not permitted by statutory regulation, users will need to obtain permission from the license holder to duplicate, adapt or reproduce the material.

## Code Appendix

The code used in this case study is available from the GitHub repository accompanying this book: <https://github.com/MIT-LCP/critical-data-book>. Further information on the code is available from this website. The following functions are provided:

- `health_forecast.m`: This function computes 100 Monte-Carlo simulations of a 5-day health forecast and displays the results.
- `IED_transition.m`: This function computes and displays the proportion of patients in each state (Intubated, Extubated, or Dead), following the transition matrix in the intervention group.
- `MCMC_solver.m`: This function computes 10,000 Monte Carlo simulations for both the control and intervention group, and computes the distribution of ventilator-free days.

## References

1. Basharin GP, Langville AN, Naumov VA (2004) The life and work of A.A. Markov. *Linear Algebra Appl* 386:3–26
2. Sonnenberg FA, Beck JR (1993) Markov models in medical decision making: a practical guide. *Med Decis Mak Int J Soc Med Decis Mak* 13(4):322–338



3. Schaefer AJ, Bailey MD, Shechter SM, Roberts MS (2005) Modeling medical treatment using Markov decision processes. In: Brandeau ML, Sainfort F, Pierskalla WP (eds) *Operations research and health care*. Springer, US, pp 593–612
4. Sutton RS, Barto AG (1998) *Reinforcement learning: an introduction*. A Bradford Book, Cambridge, Mass
5. Kreke JE (2007) Modeling disease management decisions for patients with pneumonia-related sepsis [Online]. Available: <http://d-scholarship.pitt.edu/8143/>
6. Liu JS (2004) *Monte Carlo strategies in scientific computing*. Springer, New York
7. Zucchini W, MacDonald IL (2009) *Hidden Markov models for time series: an introduction using R*. Chapman and Hall/CRC, Boca Raton (2Rev Ed edition)
8. Alagoz O, Maillart LM, Schaefer AJ, Roberts MS (2004) The optimal timing of living-donor liver transplantation. *Manag Sci* 50(10):1420–1430
9. Shechter SM, Bailey MD, Schaefer AJ, Roberts MS (2008) The optimal time to initiate HIV therapy under ordered health states. *Oper Res* 56(1):20–33
10. Maillart LM, Ivy JS, Ransom S, Diehl K (2008) Assessing dynamic breast cancer screening policies. *Oper Res* 56(6):1411–1427
11. Daniel PMG, Faissol M (2007) Timing of testing and treatment of hepatitis C and other diseases. *Inf J Comput Inf*
12. Denton BT, Kurt M, Shah ND, Bryant SC, Smith SA (2009) Optimizing the start time of statin therapy for patients with diabetes. *Med Decis Mak Int J Soc Med Decis Mak* 29(3):351–367
13. Raffa JD, Dubin JA (2015) Multivariate longitudinal data analysis with mixed effects hidden Markov models. *Biometrics* 71(3):821–831
14. Scott SL, James GM, Sugar CA (2005) Hidden Markov models for longitudinal comparisons. *J Am Stat Assoc* 100:359–369
15. Srikanth P (2015) Using Markov chains to predict the natural progression of diabetic retinopathy. *Int J Ophthalmol* 8(1):132–137
16. Alterovitz R, Branicky M, Goldberg K (2008) Motion planning under uncertainty for image-guided medical needle steering. *Int J Robot Res* 27(11–12):1361–1374
17. Morris S, Devlin N, Parkin D, Spencer A (2012) *Economic analysis in healthcare*, 2nd edn. Wiley, Chichester
18. Nord E, Daniels N, Kamlet M (2009) QALYs: some challenges. *Value Health* 12(Supplement 1):S10–S15
19. Torrance GW (1986) Measurement of health state utilities for economic appraisal. *J Health Econ* 5(1):1–30
20. Drummond M, Sculpher M (2005) Common methodological flaws in economic evaluations. *Med Care* 43(7 Suppl):5–14
21. Girard TD, Kress JP, Fuchs BD, Thomason JWW, Schweickert WD, Pun BT, Taichman DB, Dunn JG, Pohlman AS, Kinniry PA, Jackson JC, Canonico AE, Light RW, Shintani AK, Thompson JL, Gordon SM, Hall JB, Dittus RS, Bernard GR, Ely EW (2008) Efficacy and safety of a paired sedation and ventilator weaning protocol for mechanically ventilated patients in intensive care (awakening and breathing controlled trial): a randomised controlled trial. *Lancet Lond Engl* 371(9607):126–134
22. Roberts DJ, Haroon B, Hall RI (2012) Sedation for critically ill or injured adults in the intensive care unit: a shifting paradigm. *Drugs* 72(14):1881–1916