

MARKOV MODELS IN MEDICAL DECISION MAKING

Markov Models

Used to represent stochastic processes (random processes that evolve over time)


In the field of medical decision analysis they are particularly useful when modelling the progression of chronic disease

Markov Models in Medical Decision Making


- So far, we have assumed that the outcome values (mostly life expectancies) are given
- The terminal node represents the remaining lifetime of an individual
- Another method for estimating life expectancy is the Markov model
- Such models are useful when
 - A decision problem involves risk that is continuous over time
 - When the timing of events is important
 - When important events may happen more than once

Properties of a Markov Model

- A finite number of states of health
- All events of interest are modelled as transitions from one state to another
- Each state is assigned a **utility**
- The time horizon of the analysis is divided into equal increments of time (Markov cycles)
- Length of a cycle may be one year, one month etc.
- A common representation is the *state-transition diagram*



State Transition Diagram



Properties of a Markov Model

- The set of states must be **mutually exclusive** (a patient can be in only one state at a time) and **collectively exhaustive** (the patient has to be in one of the states at all times).
- Patients (or portions of an entire cohort of patients) move through the model according to probabilities that govern how likely it is to transition from one state to another.

Constructing a Markov Model

- Choose a set of mutually exclusive health states
- Determine possible transitions between these health states
 - State transitions
 - Transition probabilities
- Determine a clinically valid cycle length

Evaluation of a Markov Process

When the only attribute of interest is **duration of survival**

- Add together the average time spent in the individual states

- Expected Utility $\sum_{s=1}^n t_s$ t_s : time spent in state s

Evaluation of a Markov Process

Incorporating quality of life

Incremental utility: The utility associated with spending one cycle in a particular state

Incremental utility of the DISABLED state =0.7

(i.e. Spending the cycle in DISABLED state contributes 0.7 quality-adjusted **cycles** to the expected utility)

– Expected Utility
$$\sum_{s=1}^n t_s u_s$$

Utility for the entire process=total number of cycles spent in each state*incremental utility of that state

t_s : time spent in state s

State	Incremental Utility	Time spent in that state
WELL	1	2.5 cycles
DISABLED	0.7	1.25 cycles

Quality adjusted life expectancy of the patient

$1 \times 2.5 + 0.7 \times 1.25 = 3.9$ quality-adjusted cycles

Markov States

When constructing a Markov model of disease progression

Define the disease in terms of different (mutually exclusive) states




Fig. 1. Illustrative Markov model of disease progression. Disease states are represented by ovals and possible transitions between those states are shown by arrows. Variable names adjacent to the arrows are the transition probabilities for the model. *Abbreviations:* effect = effectiveness of drug in terms of reducing disease progression; $tpDcm$ = transition probability to the death state caused as a direct result of the disease; $tpDn$ = transition probability to natural death; $tpProg$ = transition probability to progression.

Transition Probabilities

Markov Property: No memory

The stochastic process is fully described by

- The current state
- The transition probabilities

That is, the probability of moving out of a state is not dependent on the states the patient may have experienced before entering that state

One can get around this by

- using a combination of distinct states to model particular patient histories
- using time-dependent transition probabilities

Types of Markov Processes



Markov Chain: The transition probabilities are **constant over time**

Table 1 • P Matrix

		To		
		WELL	DISABLED	DEAD
From	WELL	0.6	0.2	0.2
	DISABLED	0	0.6	0.4
	DEAD	0	0	1

Time dependent Markov Processes: Transition probabilities **vary over time**

Time dependent Markov Processes



The risk of death from the disease itself is assumed to be constant
Suppose the chance of dying from the progressive form of the disease increases over time
Then one may use tunnel states (a series of temporary states that must be visited in a fixed sequence)

Transition Probabilities

Probability versus rate

Rate: Instantaneous likelihood of transition at any point in time

Probability: The proportion of a population at risk that makes a transition over time

The probability of an event that occurs at a constant rate (r) in a specified time (t)

$$p=1-e^{-rt}$$



Changing the cycle length

When changing the cycle length from, say, 1 year to 1 month one cannot simply divide the probabilities by 12!


Instead, divide the rate by 12 and then calculate the corresponding monthly transition probabilities

Use of Markov models in DA

- Models prognosis for a given patient
- It is analogous to a utility in an ordinary decision tree
- A naive approach: use a Markov process to calculate survival for a terminal node
- A better approach: Cycle tree representation



Cycle tree representation



Cancer Progression Model

Three states: Local Cancer, Metastases, Dead

Local Cancer:

Annual Mortality=2% $P(\text{Death} | \text{Local Cancer})$

Annual progression to metastases =15% $P(\text{Metastases} | \text{Local Cancer})$

Annual Cost=20K

Annual Effectiveness= 0.95 QALYs

Metastases state

Annual Mortality=10%

Annual Cost= \$50K

Annual Effectiveness=0.90 QALYs

Dead state

No cost or effectiveness

Calculation of Outcomes

Methods

- The Fundamental Matrix Solution
- Markov Cohort Simulation
- Monte Carlo Simulation

Markov cohort simulation

- A more intuitive representation
- Suppose that you want to determine the area under a curve:
 - You can solve the integral of the function describing the curve
 - You can divide the area into blocks and sum their areas

Markov cohort simulation



WELL




SICK




DEAD

INITIAL


Assume a hypothetical cohort of patients beginning the process with some distribution among the starting state (e.g. All well, 80% well and 20% disabled)



WELL



SICK



DEAD

INTERMEDIATE

For each cycle the fraction of cohort initially at each state is partitioned among all states according to the transition probabilities


This results in a new distribution of the cohort



WELL



SICK



DEAD

FINAL

The utility accrued for the cycle

$$\text{Cycle sum} = \sum_{s=1}^n f_s \times U_s$$

Fraction of the cohort at state s Incremental utility of state s

Markov cohort simulation

Table 1 • P Matrix

		To		
		WELL	DISABLED	DEAD
From	WELL	0.6	0.2	0.2
	DISABLED	0	0.6	0.4
	DEAD	0	0	1

Incremental utility of being in the disabled state=0.7

Table 2 • Markov Cohort Simulation

Cycle	WELL	DISABLED	DEAD	Cycle Sum	Cumulative Utility
Start	10,000	0	0	—	—
1	6,000	2,000	2,000	7,400	7,400
2	3,600	2,400	4,000	5,280	12,680
•	•	•	•	•	•
•	•	•	•	•	•
23	0	1	9,999	7	23,752
24	0	0	10,000	<1	23,752
Total	15,000	12,500		23,752	23,752

Markov cohort simulation

Starting distribution

A hypothetical cohort of 10000 patients begins in the WELL state

This could be some other distribution!

e.g.

WELL	DISABLED	DEAD
8000	2000	0

Table 2 • Markov Cohort Simulation

Cycle	WELL	DISABLED	DEAD	Cycle Sum	Cumulative Utility
Start	10,000	0	0	—	—
1	6,000	2,000	2,000	7,400	7,400
2	3,600	2,400	4,000	5,280	12,680
•	•	•	•	•	•
•	•	•	•	•	•
23	0	1	9,999	7	23,752
24	0	0	10,000	<1	23,752
Total	15,000	12,500		23,752	23,752

Sonnenberg FA, Beck JR. Markov models in medical decision making: a practical guide. Medical Decision Making 1993, 13(4):322-38

Markov cohort simulation

Distribution at the end of the first cycle

In accordance with the transition probabilities

$0.2 \times 10000 = 2000$ patients moved from WELL state to the DISABLED state
Similarly, 2000 moved to the DEAD state

Table 2 • Markov Cohort Simulation

Cycle	WELL	DISABLED	DEAD	Cycle Sum	Cumulative Utility
Start	10,000	0	0	—	—
1	6,000	2,000	2,000	7,400	7,400
2	3,600	2,400	4,000	5,280	12,680
.	.	.	.	.	.
.	.	.	.	.	.
23	0	1	9,999	7	23,752
24	0	0	10,000	<1	23,752
Total	15,000	12,500		23,752	23,752

Table 1 • P Matrix

		To		
		WELL	DISABLED	DEAD
From	WELL	0.6	0.2	0.2
	DISABLED	0	0.6	0.4
	DEAD	0	0	1

Markov cohort simulation

Table 2 • Markov Cohort Simulation

Cycle	WELL	DISABLED	DEAD	Cycle Sum	Cumulative Utility
Start	10,000	0	0	—	—
1	6,000	2,000	2,000	7,400	7,400
2	3,600	2,400	4,000	5,280	12,680
.	.	.	.	.	.
.	.	.	.	.	.
23	0	1	9,999	7	23,752
24	0	0	10,000	<1	23,752
Total	15,000	12,500		23,752	23,752

Cycle sum: Sum of the number of cohort members in each state multiplied by the incremental utility for that state

State	Incremental Utility
WELL	1
DISABLED	0.7
DEAD	0

For example

Cycle Sum during Cycle 2: $3600 \times 1 + 2400 \times 0.7 + 2000 \times 0 = 5280$

Markov cohort simulation

Table 2 • Markov Cohort Simulation

Cycle	WELL	DISABLED	DEAD	Cycle Sum	Cumulative Utility
Start	10,000	0	0	—	—
1	6,000	2,000	2,000	7,400	7,400
2	3,600	2,400	4,000	5,280	12,680
.	.	.	.	.	.
.	.	.	.	.	.
23	0	1	9,999	7	23,752
24	0	0	10,000	<1	23,752
Total	15,000	12,500		23,752	23,752

More and more of the cohort ends up in the DEAD state

The fraction of the cohort in the DEAD state actually is always less than 100% because during each cycle there is a finite probability of a patient's remaining alive

Markov cohort simulation

Table 2 • Markov Cohort Simulation

Cycle	WELL	DISABLED	DEAD	Cycle Sum	Cumulative Utility
Start	10,000	0	0	—	—
1	6,000	2,000	2,000	7,400	7,400
2	3,600	2,400	4,000	5,280	12,680
.	.	.	.	.	.
.	.	.	.	.	.
23	0	1	9,999	7	23,752
24	0	0	10,000	<1	23,752
Total	15,000	12,500		23,752	23,752

Therefore, the simulation is stopped:

- When the cycle cum falls below some arbitrarily small threshold (In this example, it is 1 person-cycle) or,
- When the fraction of the cohort remaining alive falls below a certain amount

Markov cohort simulation

Table 2 • Markov Cohort Simulation

Cycle	WELL	DISABLED	DEAD	Cycle Sum	Cumulative Utility
Start	10,000	0	0	—	—
1	6,000	2,000	2,000	7,400	7,400
2	3,600	2,400	4,000	5,280	12,680
.	.	.	.	.	.
.	.	.	.	.	.
23	0	1	9,999	7	23,752
24	0	0	10,000	<1	23,752
Total	15,000	12,500		23,752	23,752

The Expected Utility for this Markov Cohort Simulation

= Cumulative Utility when the cohort has been completely absorbed

Original size of the cohort

= $23,752 / 10,000 = 2.3752$ Quality adjusted cycles

Unadjusted life expectancy

= sum of the entries in the columns of WELL and DISABLED / Cohort size

= $27,500 / 10,000 = 2.75$ cycles

Adjustments to the cost and Outcome quantities

- Discounting
- Define separate variables for discount rates of costs and outcomes

Example

$$U_t = U_0 / (1+d)^t$$

The Half-Cycle Correction

- Markov model assumes that during a single cycle, each patient undergoes no more than one state transition
- In reality, patients will be moving between the different states (phases of their disease) continuously, not at discrete points in time
- Counting the membership at the beginning or at the end of a cycle will lead to errors
- Performing a Markov Simulation is analogous to calculating expected survival (area under a survival curve)
- Each rectangle under the curve is the accounting of the cohort membership during one cycle when the count is performed **at the end of each cycle**




FIGURE 8. Counting cohort membership at the end of each cycle.




FIGURE 8. Counting cohort membership at the end of each cycle.

Counting at the end of each cycle **underestimates** the survival




FIGURE 9. Counting cohort membership at the beginning of each cycle.

Counting at the beginning of each cycle **overestimates** the survival




FIGURE 10. Illustration of the half-cycle correction.

To more accurately represent the continuous nature of the state transitions we assume that the transitions occur on average, halfway through each cycle

No way to determine state membership in the middle of a cycle




FIGURE 10. Illustration of the half-cycle correction.

Consider the count at the end of each cycle to be in the middle of a cycle that begins halfway through the previous cycle and ends halfway through the subsequent cycle

Then the under and over estimations will be balanced

Shift all cycles one half cycle to the right

Add a half cycle for the starting membership at the beginning to compensate for this shift to the right

In our example:

The Expected Utility was 2.3752 Quality adjusted cycles

After half cycle correction=2.875 QA cycles

Unadjusted life expectancy was 2.75 cycles

After half cycle correction=3.25 cycles

This shift to the right makes no difference at the end of the simulation if the cohort is completely absorbed , i.e. all are dead.

If the simulation is terminated before the absorption of the cohort, the shift will result in overestimation of the expected survival.

If the cycle length is very **short** relative to average survival, the difference between the actual survival and the relative survival will be **small**.

If the cycle length is larger relative to average survival, the difference will be more significant.


Transition Probability Matrix

	Well	Sick	Dead
Well	0.75	0.20	0.05
Sick	0	0.70	0.30
Dead	0	0	1


Starting Probability Vector: [1 0 0]

Markov cohort simulation


Time




Markov cohort simulation




Quality of Life




Markov cohort simulation



Markov cohort simulation



Markov cohort simulation



Markov cohort simulation

Cycle	Well	Sick	Dead	Cycle Reward	Total Reward
0	1	0	0	0.5	0.5
1	0.75	0.20	0.05	0.85	1.35
2	0.56	0.29	0.15	0.71	2.06
3	0.42	0.32	0.26	0.58	2.64
4	0.32	0.31	0.38	0.47	3.11

Markov cohort simulation

To sum up,

- Large number of patients are followed as a cohort
- Time dependent probabilities can be easily incorporated into the analysis (cycle allows time dependence)
- Does not provide information on the distribution or variance of expected value

Monte Carlo Simulation

Determines the prognoses of a large number of individual patients

Each patient runs through model until death-repeat large number of times ($\sim 10^4$)


Provides a distribution of survival

Standard deviation

Variance



Monte Carlo Simulation



Probabilistic Sensitivity Analysis (2nd order Monte Carlo)



- Decision tree estimates and probabilities are replaced with probability distributions (normal, uniform etc.)
- The tree is evaluated many times with random values selected from each probability distribution
- Results include means and standard deviations of the expected values of each strategy

Characteristics of Markov Approaches

Feature	Markov Cohort	Monte Carlo	Matrix
Transition Probabilities	Time-dependent	Time-dependent	Constant
Incremental Utilities	Time-dependent	Time-dependent	Constant
Accuracy	Dependent on cycle length	Dependent on number of trials	Cycle-independent
Computation	Moderate	Most	Least
Variability Measures	No	Yes	Yes

Example

The decision problem is whether to implement a new drug therapy that can reduce the rate of disease progression



Calculate ICER for the drug therapy option

Estimate the corresponding cost and effect (utility) values using the Markov model

Table II. Parameter values for the illustrative Markov model of disease progression


Name	Value	Description
Transition probabilities		
tpProg	0.01	Coefficient of increase for probability of entering the progressive disease state
tpDcm	0.15	Probability of dying from the disease in a single cycle
tpDn	0.0138	Other cause mortality for age 55 to 64
	0.0379	Other cause mortality for age 65 to 74
	0.0912	Other cause mortality for age 75 to 84
	0.1958	Other cause mortality for age 85 and over
Costs (£)		
cAsymp	500	Cost of 1 cycle in the asymptomatic disease state
cProg	3000	Cost of 1 cycle in the progressive disease state
cDrug	1000	Cost of drug for 1 cycle
cDeath	1000	Cost associated with transition to the dead state
Quality-of-life adjustments		
uAsymp	0.95	Quality-of-life weight for 1 cycle in the asymptomatic disease state
uProg	0.75	Quality-of-life weight for 1 cycle in the progressive disease state
Other parameters		
Cycle	1	Length in years of 1 cycle
Effect (%)	50	Effectiveness of drug in terms of reducing disease progression
Initial age	55	The initial age (in years) at which patients are deemed to start the model
oDR (%)	6	Discount rate for outcomes
cDR (%)	6	Discount rate for costs
<i>Abbreviation: £ = pounds sterling.</i>		

Cohort Simulation


Table III. Cohort simulation for the illustrative model (no drug therapy)

Cycle	Disease state			Total
	asymptomatic	progressive	death	
0	1000	0	0	1000
1	976	10	14	1000
2	943	28	29	1000
3	902	52	46	1000
4	854	79	67	1000
5	799	109	92	1000
6	740	139	121	1000
7	678	168	154	1000
8	614	195	191	1000
9	551	218	231	1000
10	488	237	275	1000

Cumulative proportion of the cohort by disease state over time



Without drug therapy



With drug therapy

Table IV. Results of a 1000-patient cohort simulation for the illustrative model

Strategy	Cost (£)	QALYs
No drug	9 264 897	7756
Drug	16 155 440	8625
Difference	6 890 544	869
ICER	7931 per QALY	

Abbreviations: ICER = incremental cost-effectiveness ratio; QALYs = quality-adjusted life-years; £ = pounds sterling.

Monte Carlo Simulation (Individual Simulation)

Example:

- First patient: 15 years in the asymptomatic phase and a further 5 years in the progressive state before dying
- As an individual, they are predicted to cost £10 423 while they accrue 10.5 QALYs
- Second patient: 5 years in the asymptomatic phase and only one year in the progressive state before dying
- Has a predicted cost of £4886 and will have accrued 4.5 QALYs
- Average these costs and effects over a large number of patients

Monte Carlo Simulation (Individual Simulation)

Table V. Results of 10 000 simulations of individual patients for the illustrative model

Strategy	Cost (SD) [£]	QALYs (SD)
No drug	9258 (5658)	7.74 (2.89)
Drug	16 108 (5977)	8.59 (3.09)
Difference	6850	0.85
ICER	8059 per QALY	

Abbreviations: ICER = incremental cost-effectiveness ratio; QALYs = quality-adjusted life-years; SD = standard deviation; £ = pounds sterling.

Table IV. Results of a 1000-patient cohort simulation for the illustrative model

Strategy	Cost (£)	QALYs
No drug	9 264 897	7756
Drug	16 155 440	8625
Difference	6 890 544	869
ICER	7931 per QALY	

Abbreviations: ICER = incremental cost-effectiveness ratio; QALYs = quality-adjusted life-years; £ = pounds sterling.

Gives an exact solution for the chosen cycle length

Table V. Results of 10 000 simulations of individual patients for the illustrative model

Strategy	Cost (SD) (£)	QALYs (SD)
No drug	9258 (5658)	7.74 (2.89)
Drug	16 108 (5977)	8.59 (3.09)
Difference	6850	0.85
ICER	8059 per QALY	

Abbreviations: ICER = incremental cost-effectiveness ratio; QALYs = quality-adjusted life-years; SD = standard deviation; £ = pounds sterling.

Gives an estimate of the likely variance associated with the parameters estimated by the model

“First order” Monte Carlo simulation

- It is also possible to allow the parameters (e.g. cost and probability) vary over a given range with a given distribution
- This is called second order Monte Carlo Simulation (more on this later)