Clarivate™

Clarivate Consulting Services

Building Better Models: A Simulation Case Study Demonstrating the Crucial Role of Ergodicity in Building Robust Markov Chains

Cornell M[a], Baines D[b], Howells R[c]

[a]Clarivate Evidence, Value, & Access Consulting, Health Economics, Philadelphia, US; [b]Clarivate Evidence, Value, & Access Consulting, Health Economics, London, UK; [c]Clarivate Evidence, Value, & Access Consulting, Health Economics, Manchester, UK

Code: EE482

Introduction

Markov Chains

Health economists regularly utilise MMs to help construct Cost-Effectiveness Models (CEMs) for a whole range of diseases and conditions, including prevention programmes, diagnostics, policy, treatments, and rehabilitation. Among the four types, Markov Chains (MCs) have emerged as the most popular choice of MM used by health economists building CEMs, with Hidden Markov models being used on rare occasions. In this paper, we use MC to represent the general concept and mathematical formulation of the memory-less stochastic process, while Markov Chain Model (MCM) specifically refers to the real-world application of the approach.

Health economists commonly use MCMs to simulate the dynamic progression of patient cohorts described in terms of connected health states. This information is then used to generate cost and benefit data for inclusion in CEMs and sophisticated BIMs.

- Formally, an n-state MC can be defined as a tuple (*S*, *P*), where:
- *S* is the set of possible states that the system can be in.
- *P* is the transition probability matrix is an $n \cdot n$ matrix, where $p_{i,j}$ represents the probability of transitioning from state s_i to state s_j in one step.

$$P = \begin{bmatrix} p_{1,1} & \cdots & p_{1,n} \\ \vdots & \ddots & \vdots \\ p_{n,1} & \cdots & p_{n,n} \end{bmatrix}$$

- P must satisfy the following conditions:
  - $P_{i,j} \geq 0$: Transition probabilities are non-negative.
  - $\sum_{j=1}^{n} p_{i,j} \,|\forall j \in S = 1$: The sum of probabilities for all possible transitions from state *i* is 1.

During simulations, MCMs cycle through identified sets of states, *S*, guided by associated matrices of probabilities, *P*. This results in stochastic processes that describe the sequence of states that systems go through as they change. In an MCM, the analysis starts with knowledge of *S* and *P* only. The particular values that a random MCM will take are, a priori, unknown. Symbolically, random variables are usually denoted using uppercase letters, while the particular values that they realise are usually denoted by lowercase letters. For instance, we denote a random variable by *X* to represent the MC, with function *X(t)* representing its stochastic process. After we know the *S* and *P* values, we can calculate the probability distribution, *λ*; ergo, $\lambda_{s,t} = P(X(t) = s)$ = the probability distribution of patients in state *s* at time *t*. Additionally, the sequence of *X(0),X(1),...,X(T)* must satisfy the rule of conditional independence such that:

$$P(X(t+1) = S_{t+1} | X(t) = S_t) = P(X(t+1) = S_{t+1} | X(0) = S_0, X(1) = S_1, \ldots, X(t) = S_t)$$

This assumption of independence, also known as the Markov or memory-less property, is useful analytically because, by ignoring history, computational complexity can be reduced. This, in turn, can reduce the data required for modelling. Combined, these benefits often make MCs simpler and easier to build than many other model types, particularly if meaningful probabilities can be assigned to state transitions.

Introduction (continued)

Ergodicity

An MC is a stochastic process that evolves through a sequence of states, transitioning from one state to another based on transition probabilities. The dynamics of an MCM are completely captured by its initial distribution (that is, the probability distribution, λ, that describes the probabilities of starting the chain in each possible state at the beginning of the process) and its TPM, *P*. The initial distribution sums to 1, reflecting the certainty that the chain must start in one or more of the possible states. The starting position is crucial in determining the behaviour of MCMs, especially in the initial steps of the process. As they cycle, MCMs tend to converge to a stationary distribution regardless of their initial position. However, building MCMs with both initial distributions and transition probabilities that each sum to one does not, by itself, ensure that the stochastic process will:

- Converge on a unique distribution where all states have a > 0 distribution probability.

Understanding the difference between stationary and unique distributions is important. The former occurs when MCMs settle into a long-term distribution, but this does not necessarily mean that the distribution has a single central location. MCMs can produce binomial (or higher) distributions where more than one state becomes prominent. In contrast, unique distributions occur when stationary distributions have only one central location. Unless they reach a unique distribution, MCMs can fail to be analytically useful. For instance, Therefore, testing for ergodicity (and examining its potential causes) is a task that model-building health economists should perform before using MCMs as drivers for their CEMs. An MCM is ergodic if it satisfies two conditions:

- **Irreducibility:** An MC is irreducible if it is possible to reach any state from any other state in a finite number of steps. In other words, there should be no isolated states or subsets of states that cannot be reached from the rest of the states.
- $0 > \sum_{j=1}^{n} p_{i,j} \,|\forall i,\ j \in S$: The sum of probabilities for each state j is positive.
- **Aperiodicity:** An MC is aperiodic if it lacks regular patterns or cycles in its state transitions. Thus, the chain will eventually revisit any state, irrespective of the initial state, without following fixed or predictable patterns.

If the conditions of irreducibility and aperiodicity are satisfied, then an MCM is ergodic and will settle into a unique stationary distribution. Over time, the probabilities of being in each state will converge to the stationary distribution, regardless of the initial state. The shape of the Markov trace should show an unhindered path from convergence to steady state, which occurs because the MCM is ergodic.

Introduction

Demonstrate the importance of ergodicity in building health economics models in terms of model stability and results reliability.

Methods

Through a series of simulation MCMs in R, the importance as well as impact of ergodicity on model structure and common pitfalls were identified.

Results

Example Ergodic MCM

The table below represents the transition probability matrix for an example ergodic MCM. As it is possible to reach any state from any other state in a finite number of steps, the MCM is irreducible.

		To State		
		State A	State B	State C
From State	State A	0.7	0.2	0.1
	State B	0.3	0.4	0.3
	State C	0.2	0.5	0.3

Figure 1 below showcases the shape of the Markov trace that is an unhindered path of convergence from λ to steady state, which occurs because the MCM is aperiodic. The irreducibility and aperiodicity shown verify the ergodicity of the example MCM.

Figure 1: Markov Trace for Example Ergodic MCM

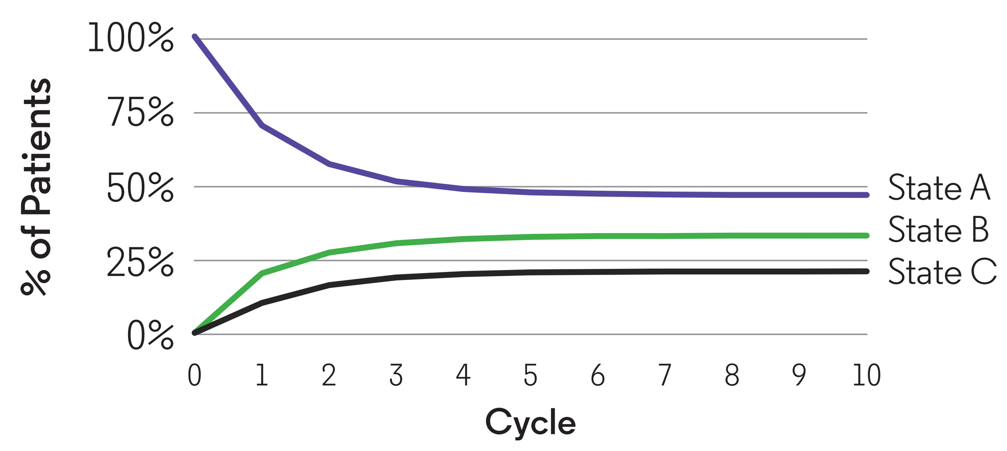

Common Pitfalls Preventing Ergodicity

1 Absorbing States / Dangling Nodes

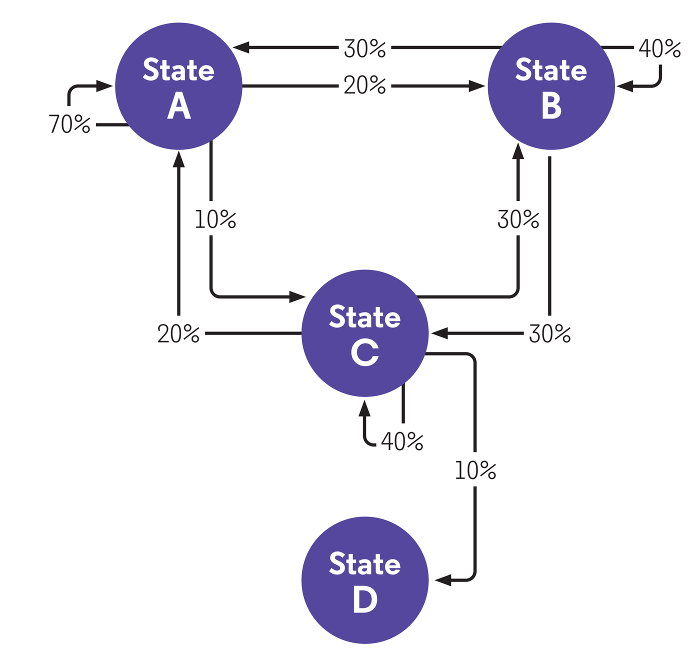

An absorbing state, also known as a dangling node, within the chain that acts as a dead end, which is characterised by transitions leading into the state from other state(s) but lacking transitions out, so the process becomes trapped, represented by the absorbing state D. Therefore, for at least one state, s ∈ S, the sum of probabilities for all possible transitions from state s is 0: $0 = \sum_{j=1}^{n} p_{s,j} \,|\forall j \in S$.

2 Artificial Linkages / Sub-Systems

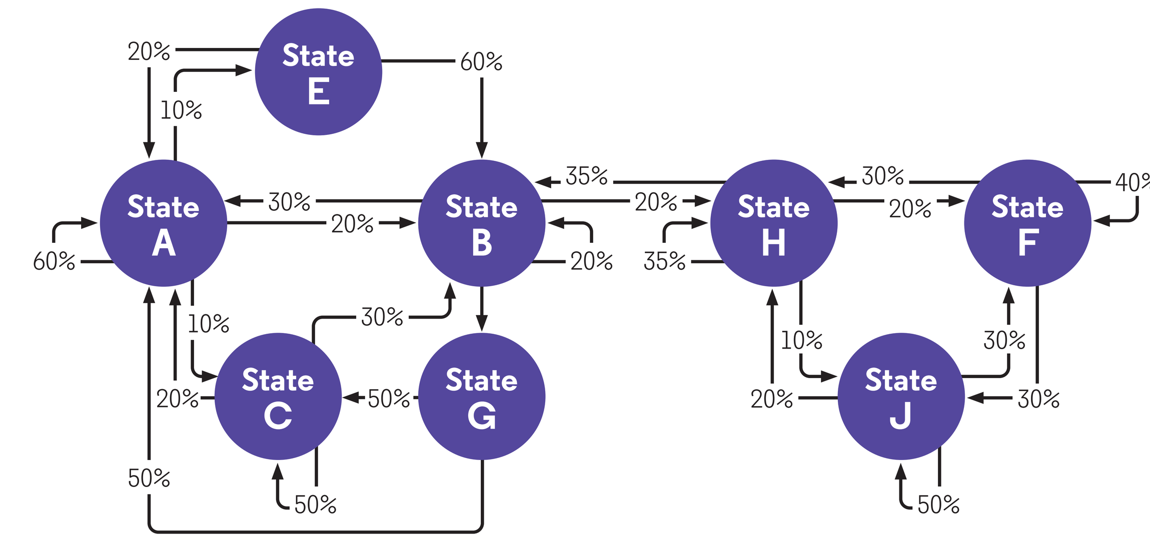

A sub-system is present as the transition connecting the main system to the sub-system (containing states H, F, & J) will affect the patterns of movement, as well as the convergence of the overall stochastic process. This may impact aperiodicity.

If the transitions between state B and state H are not direct relationships, an artificial linkage would be present, thus not observable in practice.

Validation with Non-Ergodic MCM

In scenarios where the model is non-ergodic due to an absorbing state is required, such as the specific inclusion of a death state due to the implementation of health state-specific mortality instead of an aggregate mortality rate applied to the population as a whole or in oncology modelling, it is still possible to validate would otherwise be ergodic by substituting the transition probabilities array for each state, s ∈ S, where the $0 = \sum_{j=1}^{n} p_{s,j} \,|\forall j \in S$ with the probability distribution, λ, that describes the probabilities of starting the chain in each possible state at the beginning of the process.

Conclusion

Health economists commonly employ MCMs, especially in chronic disease modelling, due to their simplicity and ease of development. Health economics analyses typically require MCMs to reach a unique equilibrium so that expected values can be estimated for each intervention modelled.

If MCMs produce binomial (or higher) distributions, average costs and benefits for each treatment will be meaningless, and ICER calculations misleading as MCMs are used to model the path from a system's starting equilibrium to its steady-state position. If they are unable to map the path of convergence to stability, MCMs may fail to meet their purpose, and their construction should be examined and questioned.

Ergo, consideration for ergodicity in model conceptualisation and execution is crucial. Even where it is not viable for the MCM to be ergodic, lessons from ergodic MMs, such as the elimination of artificial linkages and sub-systems, can be implemented to improve non-ergodic MCMs.

Presented at ISPOR 2024, Atlanta, GA, USA | May 5-8, 2024