Stochastic Models
Lecture 1 Part II
Markov Chains

Nan Chen
MSc Program in Financial Engineering
The Chinese University of Hong Kong (ShenZhen)

Sept. 13, 2017

Outline

1. State Classification: Transient and recurrence
2. Long-Term Behavior of Markov Chains

2.1 State Classification

Asymptotic Behavior of Markov Chains

- It is frequently of interest to find the asymptotic behavior of P_{ij}^n as $n \to +\infty$.
- One may expect that the influence of the initial state recedes in time and that consequently, as $n \to +\infty$, P_{ij}^n approaches a limit which is independent of i.
- In order to analyze precisely this issue, we need to introduce some principles of classifying states of a Markov chain.

Example V

- Consider a Markov chain consisting of the 4 states 0, 1, 2, 3 and having transition probability matrix

$$\begin{pmatrix} 1/2 & 1/2 & 0 & 0 \\ 1/2 & 1/2 & 0 & 0 \\ 1/4 & 1/4 & 1/4 & 1/4 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

- What is the most improbable state after 1,000 steps by your intuition?

Accessibility and Communication

- State j is said to be **accessible** from state i if for some n, $P_{ij}^{n} > 0$.
  - In the previous slide, state 3 is accessible from state 2.
  - But, state 2 is not accessible from state 3.
- Two states i and j are said to **communicate** if they are accessible to each other. We write $i \leftrightarrow j$.
  - States 0 and 1 communicate in the previous example.

Simple Properties of Communication

- The relation of communication satisfies the following three properties:
  - State i communicates with itself;
  - If state i communicates with state j, then state j communicates with state i;
  - If state i communicates with state j, and state j communicates with state k, then state i communicates with state k.

State Classes

- Two states that communicate are said to be in the same **class**.
- It is an easy consequence of the three properties in the last slide that any two classes are either identical or disjoint. In other words, the concept of communication divides the state space into a number of separate classes.
- In the previous example, we have three classes: $\{0,1\},\{2\},\{3\}$.

Example VI: Irreducible Markov Chain

- Consider the Markov chain consisting of the three states 0, 1, 2, and having transition probability matrix

$$\begin{pmatrix} 1/2 & 1/2 & 0 \\ 1/2 & 1/4 & 1/4 \\ 0 & 1/3 & 2/3 \end{pmatrix}$$

  How many classes does it contain?

- The Markov chain is said to be **irreducible** if there is only one class.

Recurrence and Transience

- Let f_i represent the probability that, starting from state i, the process will ever return to state i.
- We say a state i is **recurrent** if $f_i = 1$.
- It is easy to argue, that if a state is recurrent, then, starting from this state, the Markov chain will return to it again, and again, and again --- in fact, infinitely often.

Recurrence and Transience (Continued)

- A non-recurrent state is said to be **transient**, i.e., a transient state i satisfies $f_i < 1$.
- Starting from a transient state i,
  - The process will never again revisit the state with a positive probability $1 - f_i$;
  - The process will revisit the state just *once* with a probability $f_i(1 - f_i)$;
  - The process will revisit the state just *twice* with a probability $f_i^2(1 - f_i)$;
  - ……

Recurrence and Transience (Continued)

- From the above two definitions, we can easily see the following conclusions:
  - The number of time periods that the process will visit a transient state has a geometric distribution
  - A transient state will only be visited a finite number of times.
  - In a finite-state Markov chain not all states can be transient.
- In Example V, states 0, 1, 3 are recurrent, and state 2 is transient.

One Commonly Used Criterion of Recurrence

- **Theorem:** A state i is recurrent if and only if

$$\sum_{n=1}^{+\infty} P_{ii}^{n} = +\infty.$$

- You may refer to Example 4.18 in Ross to see one application of this criterion to prove that one-dimensional symmetric random walk is recurrent.

Recurrence as a Class Property

- **Theorem:** If state i is recurrent, and state i communicates with state j, then state j is recurrent.
- Two conclusions can be drawn from the theorem:
  - Transience is also a class property.
  - All states of a finite irreducible Markov chain are recurrent.

Example VII

- Let the Markov chain consisting of the states 0, 1, 2, 3, and having transition probability matrix

$$\begin{pmatrix} 0 & 0 & 1/2 & 1/2 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}$$

Determine which states are transient and which are recurrent.

Example VIII

- Discuss the recurrent property of a one-dimensional random walk.
- Conclusion:
  - Symmetric random walk is recurrent;
  - Asymmetric random walk is not.

2.2 Long-Run Proportions and Limiting Probabilities

Long-Run Proportions of MC

- In this section, we will study the long-term behavior of Markov chains.
- Consider a Markov chain $(X_1, X_2, \cdots, X_n, \cdots)$. Let π_j denote the long-run proportion of time that the Markov chain is in state j, i.e.,

$$\pi_j = \lim_{n \to +\infty} \frac{\#\{1 \le i \le n : X_i = j\}}{n}.$$

Long-Run Proportions of MC (Continued)

- A simple fact is that, if a state j is transient, the corresponding $\pi_j = 0$.
- Therefore we only consider recurrent states in this subsection:

  Let $N_j = \min\{k > 0 : X_k = j\}$, the number of transitions until the Markov chain makes a transition into the state j. Denote m_j to be its expectation, i.e. $m_j = E[N_j \mid X_0 = j]$.

Long-Run Proportions of MC (Continued)

- **Theorem:**

  If the Markov chain is irreducible and recurrent, then for any initial state

$$\pi_j = \frac{1}{m_j}.$$

Positive Recurrent and Null Recurrent

- **Definition:** we say a state is **positive recurrent** if $m_j < +\infty$; and say that it is **null recurrent** if $m_j = +\infty$.
- It is obvious from the previous theorem, if state j is positive recurrent, we have $\pi_j > 0$.

How to Determine π_j ?

- **Theorem:** Consider an irreducible Markov chain. If the chain is also positive recurrent, then the long-run proportions of each state are the unique solution of the equations:

$$\pi_j = \sum_i \pi_i p_{ij}, \forall j \qquad \sum_j \pi_j = 1.$$

If there is no solution of the preceding linear equations, then the chain is either transient or null recurrent and all $\pi_j = 0$.

Example IX: Rainy Days in Shenzhen

- Assume that in Shenzhen, if it rains today, then it will rain tomorrow with prob. 60%; and if it does not rain today, then it will rain tomorrow with prob. 40%. What is the average proportion of rainy days in Shenzhen?

Example IX: Rainy Days in Shenzhen (Solution)

- Modeling the problem as a Markov chain:

$$P = \begin{bmatrix} 0.6 & 0.4 \\ 0.4 & 0.6 \end{bmatrix}.$$

- Let π_0 and π_1 be the long-run proportions of rainy and no-rain days. We have

$$\pi_0 = 0.6\pi_0 + 0.4\pi_1; \qquad \pi_1 = 0.4\pi_0 + 0.6\pi_1;$$

$$\pi_0 + \pi_1 = 1.$$

$\pi_0 = \pi_1 = 1/2.$

Example XI: A Model of Class Mobility

- A problem of interest to sociologists is to determine the proportion of a society that has an upper-, middle-, and lower-class occupations.
- Let us consider the transitions between social classes of the successive generations in a family. Assume that the occupation of a child depends only on his or her parent's occupation.

Example XI: A Model of Class Mobility (Continued)

- The transition matrix of this social mobility is given by

$$P = \begin{bmatrix} 0.45 & 0.48 & 0.07 \\ 0.05 & 0.70 & 0.25 \\ 0.01 & 0.50 & 0.49 \end{bmatrix}.$$

That is, for instance, the child of a middle-class worker will attain an upper-class occupation with prob. 5%, will move down to a lower-class occupation with prob. 25%.

Example XI: A Model of Class Mobility (Continued)

- The long-run proportion π_i thus satisfy

$$\pi_0 = 0.45\pi_0 + 0.05\pi_1 + 0.01\pi_2;$$

$$\pi_1 = 0.48\pi_0 + 0.70\pi_1 + 0.50\pi_2;$$

$$\pi_2 = 0.07\pi_0 + 0.25\pi_1 + 0.49\pi_2;$$

$$1 = \pi_0 + \pi_1 + \pi_2.$$

Hence,

$$\pi_0 = 0.07, \pi_1 = 0.62, \pi_2 = 0.31.$$

Stationary Distribution of MC

- For a Markov chain, any set of $\{\pi_i\}$ satisfying

$$\pi_j = \sum_i \pi_i p_{ij}, \forall j, \text{ and } \sum_j \pi_j = 1,$$

is called a **stationary probability distribution** of the Markov chain.

- **Theorem**: If the Markov chain starts with an initial distribution $\{\pi_i\}$, i.e., $P(X_0 = i) = \pi_i, \forall i,$ then

$$P(X_t = i) = \pi_i,$$

for all state i and $t \geq 0$.

Limiting Probabilities

- Reconsider Example I (Rainy days in Shenzhen):

  Please calculate what are the probabilities that it will rain in 10 days, in 100 days, and in 1,000 days, given it does not rain today.

Limiting Probabilities (Continued)

- That transforms to the following problem:

  Given

$$P = \begin{bmatrix} 0.7 & 0.3 \\ 0.4 & 0.6 \end{bmatrix},$$

  what are P^{10}, P^{100}, and P^{1000}?

Limiting Probabilities (Continued)

- With the help of MATLAB, we have

$$P^{10} = \begin{bmatrix} 0.5714311021 & 0.4285688979 \\ 0.5714251972 & 0.4285748028 \end{bmatrix},$$

$$P^{100} = \begin{bmatrix} 0.571428571428857 & 0.428571428571427 \\ 0.571428571428857 & 0.428571428571427 \end{bmatrix},$$

$$P^{1000} = \begin{bmatrix} 0.571428571428556 & 0.428571428571417 \\ 0.571428571428556 & 0.428571428571417 \end{bmatrix}.$$

Limiting Probabilities (Continued)

- Observations:
  - As $n \to +\infty$, P_{ij}^n converges;
  - The limit of P_{ij}^n does not depend on the initial state i.
  - The limits coincide with the stationary distribution of the Markov chain

$$\pi_0 = 0.7\pi_0 + 0.3\pi_1; \quad \pi_1 = 0.4\pi_0 + 0.6\pi_1;$$
$$\pi_0 + \pi_1 = 1.$$

$$\pi_0 = 4/7 \approx 0.571428571428571$$
$$\pi_1 = 3/7 \approx 0.428571428571429$$

Limiting Probabilities and Stationary Distribution

- **Theorem:** In a positive recurrent (*aperiodic*) Markov chain, we have

$$\lim_{n\to+\infty} P_{ij}^n = \pi_j,$$

where $\{\pi_j\}$ is the stationary distribution of the chain.

In words, we say that a positive recurrent Markov chain will reach *an equilibrium/a steady state* after long-term transition.

Periodic vs. Aperiodic

- The requirement of aperiodicity in the last theorem turns out to be very essential.
- We say that a Markov chain is **periodic** if, starting from any state, it can only return to the state in a multiple of $d > 0$ steps. Otherwise, we say that it is **aperiodic.**
- Example: A Markov chain with period 3.

$$P = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}.$$

Summary

- In a positive recurrent aperiodic Markov chain, the following three concepts are equivalent:
  - Long-run proportion;
  - Stationary probability distribution;
  - Long-term limits of transition probabilities.

The Ergodic Theorem

- The following result is quite useful:

  Theorem: Let $\{X_n, n \geq 1\}$ be an irreducible Markov chain with stationary distribution $\{\pi_j\}$, and let $r(\cdot)$ be a bounded function on the state space. Then,

$$\lim_{n \to +\infty} \frac{\sum_{i=1}^{n} r(X_i)}{n} = \sum_j r(j)\pi_j.$$

Example XII: Bonus-Malus Automobile Insurance System

- In most countries of Europe and Asia, automobile insurance premium is determined by use of a Bonus-Malus (*Latin* for Good-Bad) system.
- Each policyholder is assigned a positive integer valued state and the annual premium is a function of this state.
- Lower numbered states correspond to lower number of claims in the past, and result in lower annual premiums.

Example XII (Continued)

- A policyholder's state changes from year to year in response to the number of claims that he/she made in the past year.
- Consider a hypothetical Bonus-Malus system having 4 states.

State	Annual Premium ($)
1	200
2	250
3	400
4	600

Example XII (Continued)

- According to historical data, suppose that the transition probabilities between states are given by

$$P = \begin{pmatrix} 0.6065 & 0.3033 & 0.0758 & 0.0144 \\ 0.6065 & 0 & 0.3033 & 0.0902 \\ 0 & 0.6065 & 0 & 0.3935 \\ 0 & 1 & 0.6065 & 0.3935 \end{pmatrix}.$$

- Find the average annual premium paid by a typical policyholder.

Example XII (Solution)

- By the ergodic theorem, we need to compute the corresponding stationary distribution first:

$$\pi_1 = 0.3692, \pi_2 = 0.2395, \pi_3 = 0.2103, \pi_4 = 0.1809.$$

- Therefore, the average annual premium paid is

$$200\pi_1 + 250\pi_2 + 400\pi_3 + 600\pi_4 = 326.375.$$

Homework Assignments

- Read Ross Chapter 4.3 and 4.4.
- Answer Questions:
  - Exercises 18, 20 (Page 263, Ross)
  - Exercises 23(c), 24 (Page 264, Ross)
  - Exercises 37 (Page 266, Ross)
  - Exercise 67 (page 272, Ross).

  Due on Sept. 29, Fri.