Lecture 27: Markov Chains and Martingales

This material is not covered in the textbooks. These notes are still in development. The first part of the lecture will present results that are useful, but I will make no attempt to prove them here. Still, you should know about them.

27.1 ⋆ Ergodic theorems for Markov chains

An ergodic theorem states roughly that time-series averages of a stochastic process converge to cross-sectional probabilities. If that does not make sense, bear with me.

Recall Theorem 26.9.3:

26.9.3 Theorem (Convergence to the invariant distribution) *For a Markov chain with transition matrix P, if the chain is irreducible and aperiodic, then the invariant distribution π is unique, and for any initial distribution λ, the sequence λP^n converges to π.*

In particular, for any states i and j

$$p_{ij}^{(n)} \to \pi_j \text{ as } n \to \infty.$$

The last part says that no matter which state i the chain starts in, the probability that it will be in state j after a sufficiently large number n of transition is very close to π_j. Does that mean that the long-run fraction of time the chain spends in state j is also π_j. An ergodic theorem tells you when the answer is yes.

Before proceeding, I want to go back and refine the notion of recurrence. Recall that $T_i = \inf\{t : X_t = i\}$ is called the first passage time to state i, and $\mathbb{P}_i(E) = \mathbb{P}\left(E \mid X_0 = i\right)$. Theorem 26.13.4 says that:

For a Markov chain, either

1. $\mathbb{P}_i(T_i < \infty) = 1$, in which case i is recurrent and $\sum_{n=1}^{\infty} p_{ii}^{(n)} = \infty$, or else
2. $\mathbb{P}_i(T_i < \infty) < 1$, in which case i is transient and $\sum_{n=1}^{\infty} p_{ii}^{(n)} < \infty$.

Consequently, every state is either recurrent or transient.

But there are two ways to be recurrent, depending on the expected value of the first passage time.

27.1.1 Definition *A state i is **positive recurrent** if i is recurrent and $\boldsymbol{E}_i T_i < \infty$ (where $\boldsymbol{E}_i$ denotes the expectation taken with respect to the conditional distribution $\mathbb{P}_i$). That is, the expected time to return to i is finite.*

*A state i is **null recurrent** if i is recurrent and $\boldsymbol{E}_i T_i = \infty$. That is, the expected time to return to i is infinite.*

For an irreducible chain, if one state is positive recurrent, so is every state. The following may be found in Norris [5, Theorem 1.7.7, p. 37].

27.1.2 Theorem *For an irreducible transition matrix P, the following statements are equivalent.*

1. *Every state is positive recurrent.*

2. *At least one state is positive recurrent.*

3. *The transition matrix P has an invariant distribution π.*

If these conditions hold, then for each i,

$$\boldsymbol{E}_i\, T_i = \frac{1}{\pi_i}.$$

27.1.3 Definition *We say that a Markov chain is positive recurrent if every state is positive recurrent.*

On the face of it null recurrence is an odd phenomenon. It says that if we start in state i, we will return to infinitely often with probability one, yet the expected time between returns is infinite. This can only happen for a countably infinite state space. Null recurrence is just another example of the fact that infinity is hard to comprehend.

We can now state an ergodic theorem for Markov chains. You can find this and a proof in Norris [5, Theorem 1.10.2, pp. 53–55].

27.1.4 Theorem (Ergodic Theorem) *Let P be an irreducible transition matrix, and let λ be a distribution on the state space S. Let $\{X_t\}$ be Markov(λ, P).*

For each $i \in S$ define

$$V_i(n) = \sum_{t=0}^{n-1} \mathbf{1}_{(X_t = i)},$$

the number of visits to state i that occur before time n.

Then

$$\mathbb{P}\left(\frac{V_i(n)}{n} \xrightarrow[n\to\infty]{} \frac{1}{m_i}\right) = 1,$$

where $m_i = \boldsymbol{E}\left(T_i \;\middle|\; X_0 = i\right) = \boldsymbol{E}_i\, T_i$. (The convention is that $1/\infty = 0$.)

In addition, if i is positive recurrent ($m_i < \infty$), we have

$$m_i = \frac{1}{\pi_i},$$

and if $f\colon S \to \boldsymbol{R}$ is a bounded function, then

$$\mathbb{P}\left(\frac{1}{n}\sum_{t=0}^{n-1} f(X_t) \xrightarrow[n\to\infty]{} \bar{f}\right) = 1,$$

where

$$\bar{f} = \sum_{i \in S} f(i)\pi_i,$$

where π is the unique invariant distribution for P.

So what does theorem really mean?

The random variable $V_i(n)/n$ is the fraction time through epoch n that is spent is state i. For an irreducible positive recurrent Markov chain, this time average converges to the stationary probability π_i of state i.

For an irreducible positive recurrent Markov chain, the fraction of times that state i transitions to state j,

$$\hat{p}_{ij} = \frac{\sum_{t=0}^{n-1} \mathbf{1}_{(X_t=i,\ X_{t+1}=j)}}{\sum_{t=0}^{n-1} \mathbf{1}_{(X_t=i)}}$$

converges to p_{ij} as $n \to \infty$ with probability one.

To understand this result a bit better, let's consider a few examples.

27.1.5 Example Recall the two-state Markov chain in Example 26.9.1. It has with transition matrix

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

It is irreducible, but has period 2, and has the unique invariant distribution $\pi = [1/2, 1/2]$. It is also positive recurrent.

Let's consider the time averages $V_1(n)/n$ spent in state 1. Starting in state 1, the state alternates between state 1 and state 2, so $V_1(n)/n \to 1/2 = \pi_1$. □

27.1.6 Example Recall the two-state Markov chain in with transition matrix

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

It is *not* irreducible, so the Ergodic Theorem above is not applicable, but it is aperiodic. There is no unique invariant distribution since every distribution is invariant.

The time averages $V_i(n)/n$ are either 0 or 1 depending on which state you start in. □

27.2 ⋆ Martingales

27.2.1 Definition *A* ***martingale*** *is a stochastic process* $\{X_t : t \in T\}$ *such that*

$$\boldsymbol{E}\,|X_t| < \infty \quad \forall t \in T,$$

and for every $t_1 < t_2 < \cdots < t_n < t_{n+1}$ *we have*

$$\boldsymbol{E}\left(X_{t_{n+1}} \;\middle|\; X_{t_n}, \dots, X_{t_1}\right) = X_{t_n}.$$

That is, the expectation *in the future conditioned on the current and any past values is simply the current* value.

Aside: The definition I gave is not as general as the standard definition. In the standard definition, there is also a collection $\{\mathcal{E}_t : t \in T\}$ of σ-algebras of events that such that if $s < t$, then $\mathcal{E}_s \subset \mathcal{E}_t$, that is, every event in $\mathcal{E}_s$ is also an event in $\mathcal{E}_t$. Such a collection is frequently called a **filtration**.[1] The random variables X_t are required to be **adapted** to the filtration, meaning that each event $(X_t \in [a, b])$ belongs to $\mathcal{E}_t$. Further, conditional expectations are defined with respect to the σ-algebra, a topic we shall not get into.[2] My definition restricts $\mathcal{E}_t$ to be $\sigma(\{X_s : s \leqslant t\})$.

[1] The term filtration comes from a movement by 20^{th} century French mathematicians to name mathematical objects with ordinary French words. They used the term *filtre* (meaning filter, but think of a funnel-like coffee filter) to indicate a family indexed by a particular kind of partial order.

[2] If X is a random variable and $\mathcal{E}'$ is a σ-subalgebra of $\mathcal{E}$, then $\boldsymbol{E}(X \mid \mathcal{E}')$ is a random variable adapted to $\mathcal{E}'$ such that for every event E in $\mathcal{E}'$, $\boldsymbol{E}\Big[\boldsymbol{E}(X \mid \mathcal{E}')\mathbf{1}_E\Big] = \boldsymbol{E}[X\mathbf{1}_E]$.

For example, if Y_i, $i = 0, 1, 2, \dots$ are independent mean-zero random variables, then

$$X_t = \sum_{n=0}^{t} Y_n$$

defines a martingale. Thus the Random Walk is a martingale, and so is the wealth of a gambler during a sequence of fair bets.

Another example of a martingale pops up in learning models. Let Z and Y_i, $i = 0, 1, 2, \dots$ be random variables with finite means (not necessarily independent or mean-zero). Then

$$X_n = \boldsymbol{E}(Z \mid Y_n, \dots, Y_0).$$

defines a martingale.

A **submartingale** (or **semi-martingale**) is a stochastic process $\{X_t : t \in T\}$ such that

$$\boldsymbol{E}\,|X_t| < \infty \quad \forall t \in T,$$

and for every $t_1 < t_2 < \cdots < t_n < t_{n+1}$ we have

$$\boldsymbol{E}(X_{t_{n+1}} \mid X_{t_n}, \dots, X_{t_1}) \geqslant X_{t_n}.$$

That is, the expected value in the future is greater than the current value. A **supermartingale** reverses the inequality. (You might think that since the study of these processes resulted from studying the wealth of a gambler, that the definitions ought to be reversed, but the early probabilists worked for the casino.) The term **semimartingale** refers to a process that is either a supermartingale or submartingale.

The Markov property says that the entire distribution of X_{t+s} depends on the past only through X_t.

In a martingale, only the expectation of X_{t+s} depends on the past only through X_t, but in a very special way.

27.3 ⋆ Martingale Convergence Theorem

One of the most important result on martingales is this.

27.3.1 Martingale Convergence Theorem (Cf. Doob [2, Theorem 4.1, p. 319].) *Let $\{X_n : n = 0, 1, 2, \dots,\}$ be a martingale. If $\lim_{n\to\infty} \boldsymbol{E}\,|X_n| = M < \infty$, then there is a random variable X_∞ with $\boldsymbol{E}\,|X_\infty| \leqslant M$ such that*

$$X_n \xrightarrow[n\to\infty]{a.s.} X_\infty.$$

Moreover, if $X_n \geqslant 0$ for all n, or if $X_n \leqslant 0$ for all n, then $M < \infty$ is satisfied, so the conclusion above follows.

Finally, if for some $q > 1$, we have $\lim_{n\to\infty} \boldsymbol{E}\,|X_n|^q < \infty$, then we may append ∞ to T, so that $\{X_n : n = 1, 2, \dots, \infty\}$ is a martingale, $\boldsymbol{E}\,|X_\infty|^q < \infty$, and $X_n \xrightarrow[n\to\infty]{q} X_\infty$.

27.3.1 On the terminology

If you look up the word martingale in a dictionary, you will find that it may come from the Portuguese *martengau*, meaning in inhabitant of Martigues (in Provence), or perhaps it comes from French via Spanish from the Arabic *al mirta‘ah*, meaning rein or check.

The first definition should refer to some kind of harness. Later definitions refer to a betting system. For example, my office dictionary [4] gives definition 1 as "A strap fastened to a horse's girth, passing between his forelegs, and fastened to the bit, or now more commonly ending in two rings, through which the reins pass. It is intended to hold down the head of the horse, and prevent him from rearing." It gives definition 3 as "Any system of betting which, in a series of bets, determines the amount to be wagered after each win or loss. The term is usually applied to dividing a specified amount desired to be won at one session into smaller unequal parts, arranged in a vertical column. By adding together the top and bottom figures after a loss, canceling them after a win, when all are crossed off, the desired amount is gained."

J. M. Hammersley [3], in a paper on a multidimensional generalization of martingales, offers this explanation of why the equestrian terminology may have been used to describe a stochastic process:

> The idea behind the terminology is the following. In gaming, a martingale is a fair gambling system, and this is probably the immediate source of the stochastic sense of "martingale." But in turn, the gaming term seems to have its origin in the equestrian sense of the word "martingale." In that sense, a martingale is a strap that prevents a horse from throwing up his head. If the horse is proceeding in the positive sense of the parameter i, and his mouth and breast are at heights y_i and z_i, respectively, above the ground at time i, then he will be moving in a steady horizontal fashion when his breast is now at the same height as his mouth was at the previous moment; thus, $y_i = z_i = y_{i-1}$ in conformity with (2.1). Since the strap checks upward but not downward movements of the head, it comes closer to what a mathematician would call a submartingale. If there are constraints from several different directions, as in (2.4), we may imagine them caused by several different straps, or by a harness.
>
> The history of martingales goes back a long way, and there are elaborate reliefs in the British Museum depicting martingales in the reigns of Tiglath-Pileser III (745-727 B.C.), Sennacherib (705-681 B.C.), and Assurbanipal (668-626 B.C.). Anderson [1] writes of the Assyrians: "They manage their horses with bit and bridle, and later reliefs show a remarkable anticipation of the modern martingale (not used as far as I know by any other ancient people). The reins are attached to a large tassel hanging below the horse's neck, which continues to provide a certain check on the horse's mouth. The rider is thus enabled to use both hands for his weapons, and can shoot the bow at full gallop." Müseler [8] and Hitchcock [4] give information about the various types of modern martingale (the standing, running, and Irish martingales), and the latter author has a colorful passage in which he says: "The standing martingale, which is used as a check to prevent the horse from throwing up his head and hitting the rider in the face, or carrying it too high, is a good remedy for stargazing, or for horses which have ewe-necks. ... This type of martingale is used universally on the polo ground."

[Note: The reference numbers in the above passage refer to Hammersley's bibliography, not this one.]

27.4 ⋆ Bayesian updated beliefs as a martingale

Here is the experiment, which is a model of Bayesian statistical inference. There is a finite set Θ of urns. Each urn has a number of balls of different colors. The (finite) set of colors is denoted $\mathcal{X}$. The probability of color x in urn θ is denoted $p_\theta(x)$. At stage 0, an urn is chosen at random according to the **prior probability** distribution P_0 on Θ. At each time t, a sample X_t is drawn at random the urn and then replaced. See Figure 27.1 for a partial tree diagram for the case where $\Theta = \{1,2,3\}$ and $\mathcal{X} = \{B, W\}$. Observing the color X_1 tells me something about the urn from which it is drawn. I can use this information to recompute the probability that the urn is urn θ_0 by using Bayes's Law: the posterior

$$P(\theta_0 \mid x) = \frac{p_{\theta_0}(x) P_0(\theta_0)}{\sum_{\theta \in \Theta} p_\theta(x) P_0(\theta)}.$$

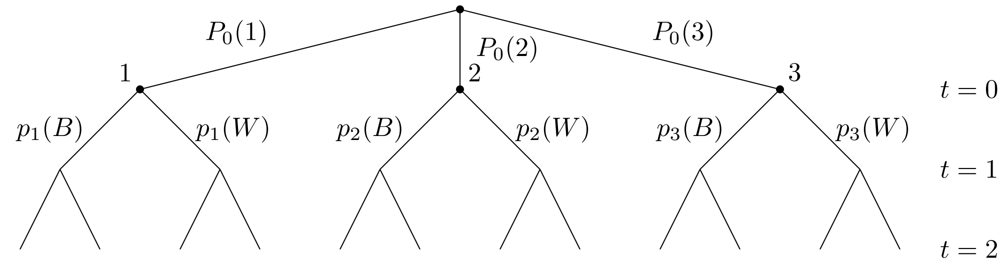

Figure 27.1. Partial tree diagram for updating urns.

(See Theorem 4.5.4.) This new probability on Θ is called the **posterior probability** on Θ given the sample x. For each $\theta \in \Theta$, define the random variable

$$\tilde{P}_1(\theta) = P(\theta \mid X_1).$$

The question is, what is the expected the expected value of the posterior probability

$$\boldsymbol{E}\,\tilde{P}_1(\theta) = \boldsymbol{E}\,P(\theta \mid X_1)?$$

Now

$$\begin{aligned}
\boldsymbol{E}\,\tilde{P}_1(\theta_0) = \boldsymbol{E}\,P(\theta_0 \mid X_1) &= \sum_{x \in \mathcal{X}} P(\theta_0 \mid X_1 = x)\,P(X_1 = x) \\
&= \sum_{x \in \mathcal{X}} \left(\frac{p_{\theta_0}(x) P_0(\theta_0)}{\sum_{\theta \in \Theta} p_\theta(x) P_0(\theta)} \right) \left(\sum_{\theta \in \Theta} p_\theta(x) P_0(\theta) \right) \\
&= \sum_{x \in \mathcal{X}} p_{\theta_0}(x) P_0(\theta_0) \\
&= P_0(\theta_0).
\end{aligned}$$

27.4.1 Remark Note that I do not say that my posterior probability will be the same as my prior probability. I am saying that my posterior probability is a random variable that depends on X_1, but the ex ante expected value of the posterior probability on Θ is the same as my prior probability.

The same argument applies into the future. For each t, define the random variables

$$\tilde{P}_t(\theta) = P(\theta \mid X_1, X_2, \dots X_t).$$

Expand on this?

We can use the Multiplication Rule for Conditional Probabilities, Theorem 4.8.1, to conclude that the posterior probability $\tilde{P}_t(\theta_0)$ I attach to urn θ_0 as a function of t independent samples (with replacement) is a martingale:

$$\boldsymbol{E}\left(\tilde{P}_{t+s}(\theta_0) \mid \tilde{P}_t(\theta_0)\right) = \tilde{P}_t(\theta_0).$$

Since probabilities are bounded, the Martingale Convergence Theorem implies that my beliefs will converge with probability one as $t \to \infty$.

But to what do the posteriors converge? If the urn that was drawn is urn θ^*, then each X_i has distribution $p(x \mid \theta^*)$, and the X_i's are independent. Then under mild regularity conditions (akin to those for conditions of maximum likelihood estimators),

$$\tilde{P}_t \xrightarrow{\mathcal{D}} \delta_{\theta^*},$$

where δ_{θ^*} assigns probability one to θ^*. The proof of this is beyond the scope of this course, but see, e.g., Degroot [1, Chapter 10].

The arguments I have just given work with densities on an infinite Θ as well as the finite case, but require a bit more care in the arguments.

27.5 ⋆ Stopping times

Given a discrete-time stochastic process $X_1, \dots, X_n \dots$, a **stopping time** is an integer-valued random variable N such that

$$\mathbb{P}(N < \infty) = 1,$$

and

$$\text{the event } (N = k) \text{ belongs to } \sigma(X_1, X_2, \dots, X_k),$$

This mean that the indicator function $\mathbf{1}_{(N=k)}$ can be written as some function h of $X_1, \dots, X_n$. In other words, you can't "peek ahead" to decide whether to stop.

27.6 ⋆ Stopped martingales

If $Z_1, Z_2, \dots$ is a martingale and N is a stopping time for this martingale, then the **stopped martingale** is

$$\bar{Z}_n = Z_{\min\{N,n\}}.$$

27.6.1 Theorem *The stopped martingale is a martingale.*

The proof is taken from Sheldon Ross and Erol Peköz [6, Lemma 3.13, p. 88].

Proof: Given n and $\omega \in \Omega$, there are two cases:

1. $N(\omega) \geqslant n$. In this case, $\min\{N, n\} = n$, so

$$\bar{Z}_n(\omega) = Z_n(\omega), \quad \bar{Z}_{n-1}(\omega) = Z_{n-1}(\omega), \quad \text{and} \quad \mathbf{1}_{(N \geqslant n)}(\omega) = 1. \tag{1}$$

Thus,

$$\begin{aligned} \bar{Z}_n(\omega) &= Z_n(\omega) && \text{by (1)}, \\ &= Z_{n-1}(\omega) + \big(Z_n(\omega) - Z_{n-1}(\omega)\big) \\ &= \bar{Z}_{n-1}(\omega) + \mathbf{1}_{(N \geqslant n)}(\omega) \cdot \big(Z_n(\omega) - Z_{n-1}(\omega)\big) && \text{by (1)} \end{aligned}$$

2. $N(\omega) < n$. In this case, $\min\{N, n\} = N(\omega)$, so

$$\bar{Z}_n(\omega) = \bar{Z}_{n-1}(\omega) = Z_{N(\omega)}(\omega), \quad \text{and} \quad \mathbf{1}_{(N \geqslant n)}(\omega) = 0. \tag{2}$$

Thus,

$$\begin{aligned} \bar{Z}_n(\omega) &= \bar{Z}_{n-1}(\omega) && \text{by (2)}, \\ &= \bar{Z}_{n-1}(\omega) + \underbrace{\mathbf{1}_{(N \geqslant n)}(\omega)}_{=0} \cdot \big(Z_n(\omega) - Z_{n-1}(\omega)\big) && \text{by (2)} \end{aligned}$$

In either case, we have

$$\bar{Z}_n = \bar{Z}_{n-1} + \mathbf{1}_{(N \geqslant n)} \cdot (Z_n - Z_{n-1}).$$

Since conditional expectation is a positive linear operator,

$$\begin{aligned}\boldsymbol{E}[\bar{Z}_n \mid Z_1,\dots,Z_{n-1}] &= \boldsymbol{E}[\bar{Z}_{n-1} \mid Z_1,\dots,Z_{n-1}] + \boldsymbol{E}[\mathbf{1}_{(N\geqslant n)} \cdot (Z_n - Z_{n-1}) \mid Z_1,\dots,Z_{n-1}] \\ &= \bar{Z}_{n-1} + \mathbf{1}_{(N\geqslant n)} \cdot \boldsymbol{E}[Z_n - Z_{n-1} \mid Z_1,\dots,Z_{n-1}] \\ &= \bar{Z}_{n-1}.\end{aligned}$$

∎

A proof of the next result may be found in Ross and Peköz [6, Theorem 3.14, pp. 88–89].

27.6.2 Theorem (Martingale Stopping Theorem)

$$\boldsymbol{E}\,Z_N = \boldsymbol{E}\,Z_1$$

if any one of the following sufficient conditions hold:

1. $\bar{Z}_n$ *are uniformly bounded.*
2. N *is bounded.*
3. $\boldsymbol{E}\,N < \infty$ *and there is some* $M < \infty$ *such that for all* n*,*

$$\boldsymbol{E}\left(Z_{n+1} - Z_n \mid Z_n\right) < M.$$

Some of the consequences of this theorem are:

- There are no gambling "systems" that guarantee positive winnings for gamblers in a fair casino who face a maximum bet limit (Condition 3).
- If there is an upper bound on family size (Condition 2), then no parental stopping rule can account for Sen's missing women.

27.7 ⋆ The Strong Markov Property: Stopped Markov chains

The Strong Markov Property asserts that if a Markov Chain is restarted after a stopping time, the continuation is also a Markov chain. The next Theorem may be found in Norris [5, Theorem 1.4.3, p. 20], and generalizes Theorem 26.4.3.

27.7.1 Theorem (Strong Markov Property) *Let* $\{X_t : t = 0,1,2,\dots\}$ *be a Markov*(λ, P) *chain, and let* T *be a stopping time for the chain. Then conditional on* $T < \infty$ *and* $X_T = i$*, the stochastic process* $\{\tilde{X}_s : s = 0,1,\dots\}$ *defined by*

$$\tilde{X}_s = X_{T+s}$$

is Markov(δ_i, P) *and independent of* $\{X_0,\dots,X_T\}$*.*

Bibliography

[1] M. H. DeGroot. 1970. *Optimal statistical decisions*. New York: McGraw-Hill.

[2] J. L. Doob. 1953. *Stochastic processes*. New York: Wiley.

[3] J. M. Hammersley. 1967. Harnesses. In L. M. Le Cam and J. Neyman, eds., *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 3: Physical Sciences*, pages 89–117. Berkeley and Los Angeles: University of California Press. http://projecteuclid.org/euclid.bsmsp/1200513623

[4] W. A. Neilson, T. A. Knott, and P. W. Carhart, eds. 1944. *Webster's new international dictionary of the English language*, second unabridged ed. Springfield, Massachusetts: G. & C. Merriam Company.

[5] J. R. Norris. 1998. *Markov chains*. Number 2 in Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge: Cambridge University Press.

[6] S. M. Ross and E. A. Peköz. 2007. *A second course in probability*. Boston: ProbabilityBookstore.com.