Chapter 13

Properties of Markov chains

13.1 Stationary distribution in a Markov chain

A Markov chain describes the random (memoryless) transitions between a set of states. Suppose that this model fits the shopping patterns of consumers. That is, customers will choose a supermarket at random from week to week, but their choice for next week depends on where they shopped this week.

Let us assume that there are only two states in this chain: Sainsbury's and Tesco (or S and T, for short). For example, we suppose that a Sainsbury's customer this week has a 70% probability of shopping there again next week, while there is a 30% probability that the customer will swap to Tesco. Conversely, suppose that a customer shopping at Tesco this week has an 80% chance of staying with Tesco for next week, and only a 20% chance of defecting to Sainsbury's. Let us assume that this is a Markov process, so customer behaviour depends only on the present and not on the past.

Let the state S be state 0 and the state T be state 1. Then this data for probabilities can be expressed as a transition matrix where the row denotes the present state, and the column the future state. Then, for this case we have:

$$\mathrm{P} = \begin{pmatrix} 0.7 & 0.3 \\ 0.2 & 0.8 \end{pmatrix} \tag{13.1}$$

where, first-row first-column, $P_{SS} = 0.7$, denotes the probability of loyalty to Sainsbury's, and so forth.

Suppose at the present the market share of customers is in the ratio: 3:2 for $S : T$. That means, if I choose a person at random from the population of shoppers, this week, there is a 60% chance that they are a customer at S. What is the probability that, any person chosen at random next week will be a customer at S ? That is, what will be the market share next week ?

To answer this question, we 'condition on' where this person might have shopped this week. So denoting the $state$ of this person next week as X_1 and the state of the person this week as X_0, then:

$$P(X_1 = S) = P(X_1 = S|X_0 = S)P(X_0 = S) + P(X_1 = S|X_0 = T)P(X_0 = T) \quad . \tag{13.2}$$

That is,

$$P(X_1 = S) = 0.7 \times 0.6 + 0.2 \times 0.4 = 0.5 \quad . \tag{13.3}$$

Since the person is chosen at random, this means that the market share of S will go down from 60% (this week) to 50% (next week). That is the probability mass has changed. Similarly for

$$P(X_1 = T) = P(X_1 = T|X_0 = S)P(X_0 = S) + P(X_1 = T|X_0 = T)P(X_0 = T) \quad . \tag{13.4}$$

For Markov chains involving n states, the partition rule applied to the probability of being in state k takes the form:

$$\boxed{P(X_1 = k) = \sum_{j=1}^{n} P(X_1 = k|X_0 = j)P(X_0 = j)} \quad . \tag{13.5}$$

We can write this in matrix form. Letting the *row vector* $\mathsf{p}^{(1)}$ be the probability distribution after one step and $\mathsf{p}^{(0)}$ be the row vector of probability distribution before the step,

$$\mathsf{p}^{(0)} = \begin{pmatrix} \mathsf{p}_1^{(0)} & \cdots & \mathsf{p}_j^{(0)} & \cdots & \mathsf{p}_n^{(0)} \end{pmatrix} \quad .$$

Then, denoting the transition matrix as:

$$\boxed{P_{ij} \equiv P(X_1 = j|X_0 = i)} \tag{13.6}$$

the equation (13.5) can be written as:

$$p_k^{(1)} = \sum_{j=1}^{n} p_j^{(0)} P_{jk} \tag{13.7}$$

or in vector notation:

$$\boxed{\mathsf{p}^{(1)} = \mathsf{p}^{(0)}\mathsf{P}} \quad . \tag{13.8}$$

So, in this case,

$$\mathsf{p}(\text{next week}) = \begin{pmatrix} 0.6 & 0.4 \end{pmatrix} \begin{pmatrix} 0.7 & 0.3 \\ 0.2 & 0.8 \end{pmatrix} \quad . \tag{13.9}$$

That is:

$$\mathsf{p}(\text{next week}) = \begin{pmatrix} 0.5 & 0.5 \end{pmatrix} \quad . \tag{13.10}$$

In general, for a homogenous Markov chain, which is the case here, we have:

$$\mathsf{p}(\text{week} \;\; \mathsf{m}) = \mathsf{p}(\text{week} \;\; \mathsf{m}-1) \times \mathsf{P} \quad . \tag{13.11}$$

So that. iterating gives:

$$\mathsf{p}(\text{week} \;\; \mathsf{m}) = \mathsf{p}(\text{week} \;\; \mathsf{m}-2) \times \mathsf{P} \times \mathsf{P} \quad . \tag{13.12}$$

Continuing in this fashion leads to:

$$\mathsf{p}(\text{week} \;\; \mathsf{m}) = \mathsf{p}(\text{week} \;\; 0) \times \mathsf{P}^{\mathsf{m}} \quad . \tag{13.13}$$

This relates the starting distribution to the distribution m weeks into the future. For example after 2 weeks, we find that:

$$\mathsf{p}(\text{week2}) = \begin{pmatrix} 0.6 & 0.4 \end{pmatrix} \begin{pmatrix} 0.7 & 0.3 \\ 0.2 & 0.8 \end{pmatrix}^2 \quad . \tag{13.14}$$

Now since:

$$\mathsf{P}^2 = \begin{pmatrix} 0.7 & 0.3 \\ 0.2 & 0.8 \end{pmatrix}^2 = \begin{pmatrix} 0.55 & 0.45 \\ 0.30 & 0.70 \end{pmatrix} \tag{13.15}$$

this gives:

$$p(\text{week2}) = \begin{pmatrix} 0.6 & 0.4 \end{pmatrix} \begin{pmatrix} 0.55 & 0.45 \\ 0.30 & 0.70 \end{pmatrix} = \begin{pmatrix} 0.45 & 0.55 \end{pmatrix} \quad . \tag{13.16}$$

That is, after the second week, the market share has shifted even more to T against S, albeit by a smaller amount.

The question is, does this loss of customers stabilise at some point ? And, if so, what is the *equilibrium* market share ?

Let us denote this (unknown) equilibrium distribution by the vector $\boldsymbol{\pi}$. Then, by its very definition, the equilibrium state will not change form one week to the next. This implies that:

$$\boxed{\boldsymbol{\pi} = \boldsymbol{\pi} \mathrm{P}} \quad . \tag{13.17}$$

This linear (matrix) equation can be recast into a more familiar form, by converting to column vectors. Taking the transpose of both sides of equation(13.17) gives:

$$\mathrm{P}^{\mathrm{T}} \boldsymbol{\pi}^{\mathrm{T}} = \boldsymbol{\pi}^{\mathrm{T}} \quad . \tag{13.18}$$

Thus, we recognise that the state we are interested in $\boldsymbol{\pi}$ is simply the eigenvector of the matrix P^{T} with eigenvalue 1. Therefore calculating $\boldsymbol{\pi}$ is a familiar problem, however we note that the solutions *may not be unique*. That is, the eigenvalue 1 may be *degenerate* (repeated). Recall that an eigenvector only provides a *direction*, that is a relation for the ratios of the elements of the vector. The normalization is a separate condition [1] .

So we impose an additional constraint on the eigenvector:

$$\boxed{\sum_{i=1}^{n} \pi_i = 1} \quad , \tag{13.19}$$

which is simply our requirement that the total probability must add to 1.

Let's return to the shopping problem, and we can either solve (13.17) or (13.18). Let us take the route of (13.17):

$$\begin{pmatrix} \pi_S & \pi_T \end{pmatrix} = \begin{pmatrix} \pi_S & \pi_T \end{pmatrix} \begin{pmatrix} 0.7 & 0.3 \\ 0.2 & 0.8 \end{pmatrix} \quad . \tag{13.20}$$

Simply multiplying out gives two simultaneous equations:

$$\pi_S = 0.7\pi_S + 0.2\pi_T \tag{13.21}$$
$$\pi_T = 0.3\pi_S + 0.8\pi_T \tag{13.22}$$
$$\tag{13.23}$$

Both equations lead to the same result (as they should):

$$\pi_T = 1.5\pi_S \tag{13.24}$$

That is, the equilibrium state is that Tesco will have 1.5 times the number of customers of Sainsbury's.

[1] That is, if u is an eigenvector of the matrix A with eigenvalue λ, then it is not unique since au is also an eigenvector, for any constant (scalar) a. This follows from the argument: $A(au) = aAu = a\lambda u = \lambda(au)$

More precisely, since the total probability should add to 1, equation (13.19):

$$\pi_S + \pi_T = 1 \tag{13.25}$$

This *normalization* equation gives us:

$$\pi_S + 1.5\pi_S = 1 \tag{13.26}$$

therefore, $\pi_S = 0.4$ and $\pi_T = 0.6$. That is, the market share will eventually stabilise with Sainsbury's at 40% and Tesco at 60%. When there are many states, the linear algebra becomes more complex. However, a simple computer program will give the answer in this case.

We note that, in the long run, starting from any initial distribution, we finish in the equilibrium state eventually, that is

$$\boldsymbol{\pi} = \mathrm{p(start)} \lim_{n\to\infty} \mathrm{P}^n \quad . \tag{13.27}$$

This leads to the important result that, if such a unique equilibrium state does exist:

$$\boxed{\lim_{n\to\infty} \mathrm{P}^n = \begin{pmatrix} \boldsymbol{\pi} \\ \boldsymbol{\pi} \\ \vdots \\ \boldsymbol{\pi} \end{pmatrix}} \quad . \tag{13.28}$$

That is:

$$\boxed{\lim_{n\to\infty} (\mathrm{P}^n)_{ij} = \pi_j \quad \text{for} \quad \text{all} \quad i} \quad . \tag{13.29}$$

Again, we emphasise that this is true, if and only if, the equilibrium state is unique.

This can be verified, since, for the jth column of equation (13.27), we have (given an arbitrary starting vector $\mathrm{p}^{(0)}$):

$$\pi_j = \sum_{i=1}^{n} p_i^{(0)} P_{ij} = \sum_{i=1}^{n} p_i^{(0)} \pi_j = \pi_j \sum_{i=1}^{n} p_i^{(0)} = \pi_j \quad . \tag{13.30}$$

This point can be illustrated for the shopping matrix, since (to 4 decimal places):

$$\mathrm{P}^4 = \begin{pmatrix} 0.4375 & 0.5625 \\ 0.3750 & 0.6250 \end{pmatrix} \quad . \tag{13.31}$$

While,

$$\mathrm{P}^{10} \approx \begin{pmatrix} 0.4006 & 0.5994 \\ 0.3996 & 0.6004 \end{pmatrix} \quad . \tag{13.32}$$

That is, the rows are beginning to converge towards the equilibrium distribution. Taking a much large number, say $n = 100$ gives, again to 4 decimal places:

$$\mathrm{P}^{100} \approx \begin{pmatrix} 0.4000 & 0.6000 \\ 0.4000 & 0.6000 \end{pmatrix} \quad . \tag{13.33}$$

13.2 Communications

In a Markov chain the connections between states are defined by the transition matrix. This matrix contains the *one-step* transition probabilities. We say that state (j) *communicates* with state (k), denoted

$j \to k$, if the Markov chain, starting at j can reach k (in a finite number of steps). That is, given the transition matrix P, there exists a finite m such that,

$$\boxed{(\mathsf{P}^m)_{jk} > 0} \quad . \tag{13.34}$$

Conversely, one can say that (j) does *not* communicates with state (k), denoted $j \nrightarrow k$, if the Markov chain, starting at j can *never* reach k, no matter how many steps taken, that is:

$$\boxed{(\mathsf{P}^m)_{jk} = 0 \quad , \quad \text{for all} \quad m \geq 0} \quad . \tag{13.35}$$

Clearly if:

$$j \to k \quad k \to l \quad \Rightarrow j \to l \quad .$$

EXAMPLE

Consider the transition graph shown above for a Markov chain involving 4 states; $\{0, 1, 2, 3\}$.

If we examine the communications from state X = 2, we note the following.

2 communicates directly (that is by a one-step transition) with 0 and 3.

2 also communicates with 1 (indirectly) via 0 (in a two-step transition).

2 also communicates with itself, through a 3-step process.

We can summarise the communications as follows:

$$2 \to \{0, 1, 2, 3\} \quad .$$

Similarly

$$3 \to 3$$

but not with any other states.

Now since $1 \to 2$, then 1 is connected (communicates) with all the states that 2 communicates with. So:

$$1 \to \{0, 1, 2, 3\} \quad .$$

and similarly:

$$0 \to \{0, 1, 2, 3\} \quad .$$

If two states *communicate* with each other, that is, $j \to k$ and $k \to j$ then we write: $j \leftrightarrow k$. Furthermore if

$$i \to j \quad \text{and} \quad j \to k \Rightarrow i \to k \quad .$$

We can prove this by noting that: $i \to j$ implies that there exists an $m \geq 0$ such that: $(\mathsf{P}^m)_{ij} > 0$ and that there exists an $n \geq 0$ such that: $(\mathsf{P}^n)_{jk} > 0$.

It follows that:

$$(\mathsf{P}^{m+n})_{ik} = \sum_s (\mathsf{P}^m)_{is} (\mathsf{P}^n)_{sk} > 0 \quad ,$$

and thus that $i \to k$.

A pair of states that communicate with each other are in the same *equivalence class*. If all the states communicate with each other, we say the the Markov chain is *irreducible*. In general this is not the case, and the states will be grouped into subsets.

13.3 Recurrent and transient states

The states in a Markov chain are classified according to the communications of the states in the subset. Let us consider two important classifications of subsets based on the states. A state can either be *transient* or *recurrent*, but not both.

For any state j,let us consider returning to j (for the first time) after n steps (or equivalently after a *time* n): The *first return time*,R, for state j can be defined as the random variable:

$$R_j = \min n : P(X_n = j|X_0 = j) > 0 \qquad . \tag{13.36}$$

Then we define a transient state as one for which there is a finite probability of never returning. That is:

$$P(R_j < +\infty) < 1 \qquad \text{transient} \qquad . \tag{13.37}$$

A recurrent state, on the other hand, is one for which

$$P(R_j < +\infty) = 1 \qquad \text{recurrent} \qquad . \tag{13.38}$$

That is, it is certain that the state will return at some time. These definitions are exclusive, and thus, a state is either transient or recurrent.

If there is certainty of return in a finite time, then for an infinite time the returns can happen infinitely often. A recurrent state will return with certainty and infinitely often. Thus, since:

$$\lim_{n\to\infty} (\mathsf{P}^n)_{jj} > 0 \qquad , \tag{13.39}$$

$$\sum_{n=0}^{\infty} (\mathsf{P}^n)_{jj} = +\infty \qquad . \tag{13.40}$$

While for a transient state we will have:

$$\lim_{n\to\infty} (\mathsf{P}^n)_{jj} = 0 \qquad , \tag{13.41}$$

that is the transient state will only be revisited a finite number of times, and in particular, the series below converges:

$$\sum_{n=0}^{\infty} (\mathsf{P}^n)_{jj} < +\infty \qquad . \tag{13.42}$$

In the example shown above, the state 3 is absorbing (and hence recurrent), while the subset $\{0, 1, 2\}$ are transient states since there is a finite probability that, once the system leaves any of these states it has a finite probability of ending up in 3 from which there is no return.

13.3.1 Classes

If state i is recurrent and $i \leftrightarrow j$ then j is also recurrent.
If state i is transient and $i \leftrightarrow j$ then j is also transient.

These statements are complementary, and only one of these requires a proof. Given that $i \leftrightarrow j$, then there exist m, n such that $(\mathsf{P}^m)_{ij} > 0, (\mathsf{P}^n)_{ji} > 0$. Now consider the recurrence of j:

$$\left(\mathsf{P}^{n'}\right)_{jj} = \sum_{r,s} (\mathsf{P}^n)_{js} \left(\mathsf{P}^{n'-m-m}\right)_{sr} (\mathsf{P}^m)_{rj} \geq (\mathsf{P}^n)_{ji} \left(\mathsf{P}^{n'-m-n}\right)_{ii} (\mathsf{P}^m)_{ij} \quad . \tag{13.43}$$

Now, it follows that:

$$\lim_{n'\to\infty} \left(\mathsf{P}^{n'}\right)_{jj} = (\mathsf{P}^n)_{ji} \lim_{n'\to\infty} (\mathsf{P}^n)_{ii} (\mathsf{P}^m)_{ij} > 0 \quad , \tag{13.44}$$

and hence that state j is recurrent.

So transience and recurrence is a property common to a subset of states, and we use this to define our classes.

There are special sub-classes, for example *absorbing* states are recurrent, trivially. But we also have states that are recurrent in a special way - *periodic states*.

13.3.2 Periodicity

The *period* $d(j)$ of a state j is the greatest common divisor (g.c.d.) of the *return times* to j.

$$d(j) = \text{gcd}\{n : p_{jj}(n \text{ steps}) > 0\}$$

If $d(j) = 1$ that is the period is one, then j is termed an *aperiodic state*.

EXAMPLE

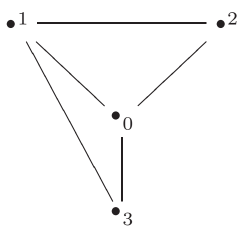

This chain has $0, 1, 2, 3$ as a (recurrent) irreducible set. This is clear from the following analysis.

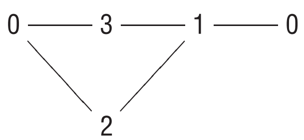

$\Rightarrow 0 \to \{0, 1, 2, 3\}$

$1 \to 0$ and 0 communicates with 0,1,2,3 $\Rightarrow 1 \to$ {0,1,2,3}
$2 \to 1$ and 1 communicates with 0,1,2,3 $\Rightarrow 2 \to$ {0,1,2,3}
$3 \to 1$ and 1 communicates with 0,1,2,3 $\Rightarrow 3 \to$ {0,1,2,3}
$\Rightarrow$ all states intercommunicate $\Rightarrow$ irreducible set/chain.

Note that 3 is recurrent (all states of an irreducible chain are recurrent)

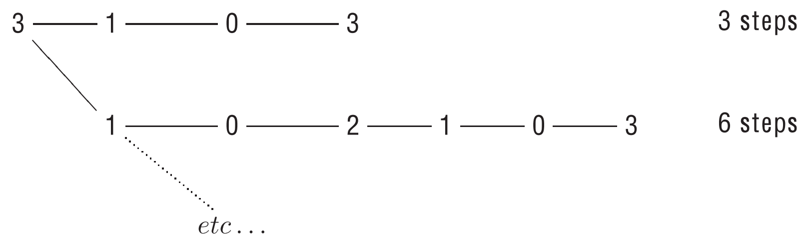

State 0 has period: $d(0) = 3$, so too $d(1) = d(2) = 3$.

EXAMPLE

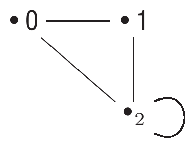

Here, state 2 is absorbing. Once 2 is reached, the system stays there forever.

Then the subset $\{0, 1, 2\}$ is *not* irreducible.

In fact, states 0 and 1 are *transient*. That is there is a finite probability that, starting at 0 (or 1) one never returns. From 0, one can go to the *absorbing state* 2 and never return.

So the set of states $\{0, 1, 2\}$ can be partitioned into two subsets:

$$\text{transient} \quad \{0,1\} \quad , \quad \text{absorbing} \quad \{2\} \quad .$$

EXAMPLE

Consider a system with 3 states represented by the (equivalent) transition graph and transition matrix shown below:

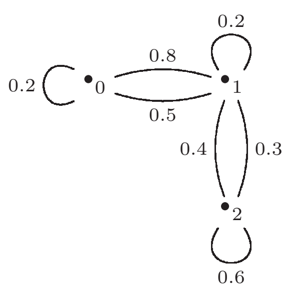

$$\Longleftrightarrow \qquad \mathbf{P} = \begin{pmatrix} 0.2 & 0.8 & 0.0 \\ 0.5 & 0.2 & 0.3 \\ 0.0 & 0.4 & 0.6 \end{pmatrix}$$

This example is an irreducible chain, every state communicates with every other state. Every state is *recurrent* and aperiodic.

13.3.3 Identification by computation

Consider the *classes* (subsets) of the following chain:

$$P = \begin{pmatrix} 0.2 & 0.8 & 0.0 & 0.0 \\ 0.5 & 0.5 & 0.0 & 0.0 \\ 0.1 & 0.2 & 0.3 & 0.4 \\ 0.0 & 0.0 & 0.0 & 1.0 \end{pmatrix} \Longleftrightarrow$$

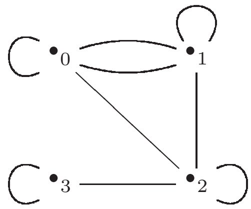

On examination, we can identify the subsets of this set of states.

$$\text{recurrent} \quad \{0,1\} \quad ; \quad \text{absorbing/recurrent} \quad \{3\} \quad ; \quad \text{transient} \quad \{2\} \quad .$$

A lazier method of analysis employs the computer to study the return probability. Recall that we spoke of transient and recurrent states in terms of the long-term transition matrix. In particular for a transient state:

$$\lim_{n\to\infty} (P^n)_{ii} = 0 \quad . \tag{13.45}$$

We can calculate the transition probabilities directly for long times, for example $n = 50$ steps. We find that:

$$P^{50} = \begin{pmatrix} 0.3846 & 0.6154 & 0 & 0 \\ 0.3846 & 0.6154 & 0 & 0 \\ 0.1648 & 0.2637 & 0.0000 & 0.5714 \\ 0 & 0 & 0 & 1.0000 \end{pmatrix} \quad .$$

Clearly, this matrix does not obey the rule (13.29) and one can conclude that a unique equilibrium state does not exist. There will be more than one equilibrium and these can be identified from this expression.

Along the diagonal of the matrix we have the return probabilities for each state corresponding to $n = 50$ steps. The fact that;

$$(P^{50})_{22} \approx 0 \quad , \quad \text{suggests} \quad \lim_{n\to\infty} (P^n)_{22} = 0 \quad .$$

That is, after a sufficiently long time (large number of steps) the probability of return tends to zero, and we have a transient state. We also note that this calculation confirms that:

$$\lim_{n\to\infty} (P^n)_{33} = 1 \quad .$$

so that state 3 is absorbing, that is $\{3\}$ is a recurrent subset. Of more interest is the subset $\{0, 1\}$ which is a recurrent subset - that is there is a finite probability of return. Moreover, their two rows are identical in this long-time limit. Indeed this is the equilibrium distribution for this subset.

This confirms the original assertion that the Markov chain is partitioned into disjoint sub-classes (subsets) of either *recurrent* or *transient* states. The numbers in the third row then represent the probabilities of ending up in states 0,1, or 3 having started from the transient state 2. These numbers are called *hitting probabilities* and will be discussed later.

Let us conclude with yet another example that displays the partitioning of subsets.

EXAMPLE

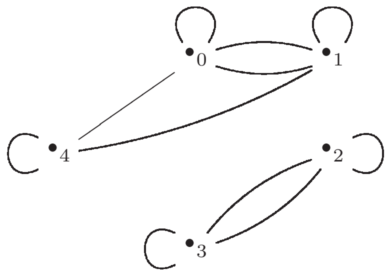

{0,1} recurrent - both 0,1 aperiodic

{2,3} recurrent - both 2,3 aperiodic

{4} is transient (finite probability of *never* returning)