DEPARTMENT OF MATHEMATICS AND STATISTICS

53304 Outline lecture notes

2. Markov Chains

We now consider processes that move around on a (usually finite) state space S in discrete time. This means that at each time point an object moves from one position in the state space to another (or it may stay at the same position). Denote the position at time t by Z_t for $t = 0, 1, \ldots$ and $Z_t \in S$.

Definition. The stochastic process Z_t for $t = 0, 1, \ldots$ is called a *Markov chain* if the following *Markov property* holds

$$\mathrm{P}(Z_{t+1} = j \mid Z_t = i, Z_{t-1} = i_{t-1}, \ldots, Z_0 = i_0) = \mathrm{P}(Z_{t+1} = j \mid Z_t = i)$$

for any $t = 0, 1, \ldots$ and $j, i_0, \ldots, i_{t-1}, i \in S$.

The Markov property is equivalent to

$$\mathrm{P}(Z_{t+k} = j \mid Z_t = i, Z_{t-1} = i_{t-1}, \ldots, Z_0 = i_0) = \mathrm{P}(Z_{t+k} = j \mid Z_t = i), \quad \forall k \geq 1.$$

Roughly this states that, given the state at time t, the behaviour after time t is independent of the behaviour before time t. In order to predict the future behaviour you need to know the current position, but information about how the process reached the current position (i.e. previous history) is of no further help.

In order to know the joint probability distribution

$$\mathrm{P}(Z_0 = i_0, Z_1 = i_1, \ldots, Z_t = i_t)$$

we introduce the *one-step transition probabilities*

$$p_{ij}^{t,t+1} = \mathrm{P}(Z_{t+1} = j \mid Z_t = i).$$

If

$$p_{ij}^{t,t+1} = p_{ij},$$

namely the one-step transition probabilities are indepdent of t (i.e. the transition probabilities are the same at all times), we say that the Markov chain is *stationary*.

We shall consider only *stationary* Markov chains. The transition probabilities can be put into a square matrix

$$P = (p_{ij}),$$

the rows and columns of which are indexed by the elements of S.

Example 1. A very simple model for the weather from day to day. If it is raining today then the probability that it will rain tomorrow is 0.8 . If it is dry today, then the probability that it will rain tomorrow is 0.4 . The state space = $\{\text{rain}, \text{dry}\}$. The transition matrix is

$$P = \begin{array}{c} \\ \text{rain} \\ \text{dry} \end{array} \begin{array}{c} \begin{array}{cc} \text{rain} & \text{dry} \end{array} \\ \begin{pmatrix} 0.8 & 0.2 \\ 0.4 & 0.6 \end{pmatrix} \end{array}.$$

Example 2. A maze used for training rats. There are 5 compartments labelled $1, \ldots, 5$ and connecting one-way doors. An untrained rat moves in a Markov chain according to the transition matrixq

$$P = \begin{pmatrix} 0 & 1/2 & 1/2 & 0 & 0 \\ 1/3 & 0 & 1/3 & 1/3 & 0 \\ 1/3 & 1/3 & 0 & 0 & 1/3 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

(Because the labels are $1, \ldots, 5$ it is not necessary to label the rows and columns of the matrix.)

Example 3. Simple *random walk* with *absorbing barriers* at ± 2 (, also known as drunkard's ruin). We shall spend some time looking in more detail at random walks later. Here the state space is $\{-2, -1, 0, 1, 2\}$. States ± 2 being absorbing barriers means that once the object reaches ± 2 it remains there. While the object is at one of the intermediate states, the probabilities of moving one step to the right/left/not moving are p, q, r, given non-negative numbers with sum 1. Here the transition matrix is

$$P = \begin{matrix} \\ -2 \\ -1 \\ 0 \\ 1 \\ 2 \end{matrix} \begin{matrix} \begin{matrix} -2 & -1 & 0 & 1 & 2 \end{matrix} \\ \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ q & r & p & 0 & 0 \\ 0 & q & r & p & 0 \\ 0 & 0 & q & r & p \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{matrix}.$$

What do you notice about the transition matrices we have examined?

In fact all the entries are non-negative and each row sum is 1. Such a (square) matrix is called a *stochastic matrix*. Note also that, if we denote by $\mathbf{1}$ the column-vector with entries labelled by the elements of S and each equal to 1, then $P\mathbf{1} = \mathbf{1}$. This can be interpreted as $\mathbf{1}$ is a right eigenvector of P corresponding to the eigenvalue 1. (Recall that in general λ is an *eigenvalue* of the square matrix P if the equation $P\mathbf{x} = \lambda\mathbf{x}$ has a solution $\mathbf{x} \neq \mathbf{0}$. Then $\mathbf{x}$ is a corresponding right eigenvector.) So stochastic matrices always have the eigenvalue 1.

Now consider what happens in two steps starting from state i.

$$\begin{aligned} p_{ij}^{(2)} &\equiv \mathrm{P}(Z_{t+2} = j \mid Z_t = i) = \sum_{k \in S} \mathrm{P}(Z_{t+2} = j \quad \& \quad Z_{t+1} = k \mid Z_t = i) \\ &= \sum \mathrm{P}(Z_{t+2} = j \mid Z_{t+1} = k \quad \& \quad Z_t = i)\mathrm{P}(Z_{t+1} = k \mid Z_t = i) \\ &\qquad\qquad \text{(defn of conditional prob., all probs conditional on } Z_t = i) \\ &= \sum \mathrm{P}(Z_{t+2} = j \mid Z_{t+1} = k)\mathrm{P}(Z_{t+1} = k \mid Z_t = i) \qquad \text{(Markov property)} \\ &= \sum_{k \in S} p_{kj} p_{ik} = [P^2]_{ij}. \qquad \text{(defn of matrix multiplication)} \end{aligned}$$

So P^2 is the transition matrix that describes movements over two time units. Similarly P^n is the matrix of 'n-step transition probabilities'.

The transition probabilities describe movements of the Markov chain from one state to another. However, this is not enough to specify the probabilistic behaviour (or law) of the process $\{Z_t\}_{t\geq 0}$. For this purpose, let us define the *initial distribution*

$$p_{i_0} = \mathrm{P}(Z_0 = i_0), \quad i_0 \in S.$$

Let us now explain that the transition matrix P and the initial distribution enable us to find, at least in principle, any probability connected with the process, such such $\mathrm{P}(Z_n = i)$ or $\mathrm{P}(Z_0 = i_0, \ldots, Z_t = i_t)$. Indeed,

$$\begin{aligned}\mathrm{P}(X_n = i) &= \sum_{k\in S} \mathrm{P}(Z_0 = k \quad \& \quad Z_n = i)\\ &= \sum_{k\in S} \mathrm{P}(Z_0 = k)\mathrm{P}(Z_n = i \mid Z_0 = k)\\ &= \sum_{k\in S} p_k p_{ki}^{(n)}.\end{aligned}$$

Moreover, compute the joint probability

$$\begin{aligned}&\mathrm{P}(Z_0 = i_0,\ Z_1 = i_1,\ \ldots, Z_{n-1} = i_{n-1},\ Z_n = i_n)\\ =&\mathrm{P}(Z_0 = i_0,\ Z_1 = i_1,\ \ldots, Z_{n-1} = i_{n-1})\\ &\times \mathrm{P}(Z_n = i_n \mid Z_0 = i_0,\ Z_1 = i_1,\ \ldots, Z_{n-1} = i_{n-1})\\ =&\mathrm{P}(Z_0 = i_0,\ Z_1 = i_1,\ \ldots, Z_{n-1} = i_{n-1})\mathrm{P}(Z_n = i_n \mid Z_{n-1} = i_{n-1})\\ =&\mathrm{P}(Z_0 = i_0,\ Z_1 = i_1,\ \ldots, Z_{n-1} = i_{n-1})p_{i_{n-1}i_n}.\end{aligned}$$

Repeating this procedure gives

$$\begin{aligned}\mathrm{P}(Z_0 = i_0,\ Z_1 = i_1,\ \ldots, Z_{n-1} = i_{n-1},\ Z_n = i_n) =&\mathrm{P}(Z_0 = i_0)p_{i_0i_1}\cdots p_{i_{n-1}i_n}\\ =&p_{i_0}p_{i_0i_1}\cdots p_{i_{n-1}i_n}.\end{aligned}$$

Similarly, the conditional probability

$$\mathrm{P}(Z_1 = i_1,\ \ldots, Z_{n-1} = i_{n-1},\ Z_n = i_n \mid Z_0 = i_0) = p_{i_0i_1}\cdots p_{i_{n-1}i_n}.$$

These show that once the initial distribution and the transition matrix are given, the probability distributions of the Markov chain $\{Z_t\}_{t\geq 0}$ are determined.

The main object of this part of the course is to study the limiting behaviour of P^n as $n \to \infty$. Sometimes there is a proper subset of the states which is *closed* in the sense that once the process enters the set it can never leave it. There may also be states that may be visited a few times, but which are eventually left for good. In many processes P^n tends to a limit. Moreover each row of P^n tends to the same limit. But there are some exceptions to this and the theory will enable us to recognise the exceptions.

Classification of states. If $p_{ii} = 1$, then i is an *absorbing state*.

A non-empty subset C of the state space is called a *closed class* if it is not possible to leave C starting from a state in C, i.e. if for all states $i \in C, j \notin C, \quad p_{ij} = 0$.

The submatrix of P defined by the rows and columns indexed by the elements of C is then a stochastic matrix.

An absorbing state on its own forms a closed class of size 1.

An *irreducible closed class* C is a closed class such that no proper subset of C is itself closed. A Markov chain is *irreducible* if S is an irreducible closed class, i.e. if there is no closed class other than S itself.

A state i is *transient* if the probability starting from i of never returning to i is positive. This occurs if there is a state j which can be reached from i in one or more steps from which it is not possible to get back to i. It can be shown that the states which are not in any irreducible closed class are all transient.

A state is *recurrent* if the probability of sooner or later returning to i starting from i is 1. If there are finitely many states, then all states in an irreducible closed class are recurrent, and to determine the closed classes and the transient states, we just need to know which elements of P are positive and which 0. (This question is more complicated in the case where there are infinitely many states. That case is not treated in this course.)

Example 1 is irreducible; in Example 2, $\{5, 6\}$ forms an irreducible closed class and states $1, \dots, 4$ are transient; in Example 3, $\{-2\}, \{2\}$ are irreducible closed classes and states $-1, 0, 1$ are transient.

Example 4. Stock control. Suppose that ordering of new stock takes place at the end of each week with a policy of not placing an order if there are one or more item remaining in stock and, when no items remain in stock, of ordering sufficient to meet any unfilled orders and make the stock up to three. The weekly demand has a known probability distribution; 0, 1, 2, or 3 with probability 0.2, 0.4, 0.3, 0.1 respectively. The demands in different weeks are independent. We shall define the state as the number of items in stock at the end of the week, counting unfilled orders as negative. The the state space is $\{-2, -1, 0, 1, 2, 3\}$. For example State -2 occurs if there was one item in stock at the end of the previous week and three items have been demanded. The transition matrix is

$$\begin{array}{c} \\ -2 \\ -1 \\ 0 \\ 1 \\ 2 \\ 3 \end{array}
\begin{array}{c} \begin{array}{cccccc} -2 & -1 & 0 & 1 & 2 & 3 \end{array} \\
\begin{pmatrix} 0 & 0 & .1 & .3 & .4 & .2 \\ 0 & 0 & .1 & .3 & .4 & .2 \\ 0 & 0 & .1 & .3 & .4 & .2 \\ .1 & .3 & .4 & .2 & 0 & 0 \\ 0 & .1 & .3 & .4 & .2 & 0 \\ 0 & 0 & .1 & .3 & .4 & .2 \end{pmatrix} \end{array} .$$

We might want to answer questions like: how often is an order placed, how often are you unable to supply an item from stock, what is the average amount of stock held?

This Markov chain is irreducible.

Example 5. Success runs. A coin is tossed independently until five heads in succession have been obtained. We denote the state here as the number of consecutive heads (up to 5) that have been obtained on the most recent tosses. S is $\{0, \ldots, 5\}$. Suppose that the probability of a head is p at each toss and denote $1 - p$ by q. How long does it take on average?

$$\begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array} \begin{array}{c} \begin{array}{cccccc} 0 & 1 & 2 & 3 & 4 & 5 \end{array} \\ \begin{pmatrix} q & p & 0 & 0 & 0 & 0 \\ q & 0 & p & 0 & 0 & 0 \\ q & 0 & 0 & p & 0 & 0 \\ q & 0 & 0 & 0 & p & 0 \\ q & 0 & 0 & 0 & 0 & p \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{array}$$

Here the state 5 is an absorbing state and the other states are transient.

Example 6. Ehrenfest urn model for the diffusion of molecules of a gas through a membrane. There are N particles in a container which has a permeable partition. At each time point one of the particles chosen at random passes through the partition. Here we record the state as the number of particles to one side of the partition, so $S = \{0, \ldots, N\}$. When the state is i, there is probability i/N that the next state is $i-1$, and probability $(N-i)/N$ that it is $i + 1$. The transition matrix is

$$\begin{array}{c} \\ 0 \\ 1 \\ 2 \\ \vdots \\ N \end{array} \begin{array}{c} \begin{array}{cccccc} 0 & 1 & 2 & 3 & \ldots & N \end{array} \\ \begin{pmatrix} 0 & 1 & 0 & 0 & \ldots & 0 \\ \frac{1}{N} & 0 & \frac{N-1}{N} & 0 & \ldots & 0 \\ 0 & \frac{2}{N} & 0 & \frac{N-2}{N} & \ldots & 0 \\ \vdots & \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & 0 & \ldots & 0 \end{pmatrix} \end{array}$$

This is an irreducible Markov chain.

This illustrates one further phenomenon which occurs — *periodicity*. In Example 6 the states that occur are altermately even and odd. If the process starts for example in state 0 then after an even number of steps it cannot avoid being in an even-numbered state, while after an odd number of steps it must be in an odd-numbered state. This complicates the description of the limiting behaviour of P^n. In fact there are two limit matrices, one for the even powers and the other for the odd powers of P.

An irreducible Markov chain is said to have *period* d if the highest common factor (hcf) of $\{n : p_{ii}^{(n)} > 0\}$ is d. It can be shown that this is the same for all states i in an irreducible Markov chain. If the period is 1 it is usual to call the Markov chain *aperiodic* (i.e. it does not have a period). Once again to recognise periodicity it is necessary only to know which entries are positive and which 0. It is possible to invent Markov chains with any period (Try it.); however any period greater than 2 will usually be obvious. Where the period is 2, the states can be put into two classes (even/odd, white/black, etc.) and all steps of the process are from one class into the other, so that the process alternates between the two classes. Sometimes you will have to think a bit to see whether this is possible.

A *stationary distribution* is a probability distribution π on S, thought of as a row-vector, such that $\pi P = \pi$. It can also be called an *equilibrium* or *steady state* distribution. Note that row i of P represents the probability distribution of the state reached in one step starting from i, and $\mathbf{y}P$ represents the probability distribution of the state reached one step after starting from a random position with probability distribution $\mathbf{y}$ (as a row-vector). The idea is that if you start the process off with an equilibrium distribution, for example by building this into a computer program simulating the process, then the probability distribution of the state at any time remains the same. It is important to realise that in doing this you do not actually look to see what the state is, because if you know it then that changes the predictions you will make about its progress. You just have to trust that the process has been started off at random according to your instructions.

Note that π is a left eigenvector corresponding to the eigenvalue 1. It follows from matrix theory that a Markov chain with finitely many states must have a left eigenvector for the eigenvalue 1. (It is not so elementary to show that this eigenvector can be take to have its entries all non-negative, so that division by a suitable constant gives a probability distribution.)

Theorem. An irreducible aperiodic Markov chain with finitely many states has a stationary distribution given by the row-vector π. Each row of P^n tends to π. If the chain is periodic, then P^n does not tend to a limit but π still represents the proportion of time that is spent in the various states in the long run.

No formal proof of this is given. However note that it is easy to show that, if P^n tends to a limit, then every row of the limit is a left eigenvector of P. For $P^{n+1} = P^n P$. Suppose that row j of the limiting matrix is $\mathbf{y}$. Then picking the jth row and letting $n \to \infty$ gives $\mathbf{y} = \mathbf{y}P$.

To investigate the limiting behaviour of P^n, first identify irreducible closed classes of states, and periodicity. Then, in each irreducible closed class, solve the equations $\mathbf{y} = \mathbf{y}P$ with the sum of coefficients equal to 1. We look at some of the examples.

Example 1 is irreducible and aperiodic. $\mathbf{y} = (2, 1)$. Normalising gives $(2, 1)/3$. This gives both rows of the limit of P^n. In the long run it rains on $2/3$ of the days.

In Example 2, the states 5,6 form an irreducible closed class, and the other states are transient. Each row of P^n tends to $(0,\ 0,\ 0,\ 0,\ 2,\ 1)/3$.

In Example 3, there are two absorbing states. We are not yet in a position to give the limit. However we can say that the three central columns of P^n all tend to 0 because the corresponding states are all transient. We shall find the probabilities of absorption at 2 starting from the intermediate states later.

Example 4 is irreducible and aperiodic. It can be shown that each row of P^n tends to $(5,\ 19,\ 40,\ 50,\ 40,\ 16)/170$. So the states $-2, -1$, which correspond to having unfilled orders, occur in the long run in $24/170$ of the weeks.

Example 5 has one absorbing state, and the other states are transient. So each row of P^n tends to (0, 0, 0, 0, 0, 1).

Example 6 is irreducible but with period 2. The case of $n = 5$ will be illustrated. It can be shown that the stationary distribution is the Binomial distribution $B(5, \frac{1}{2}) = (1, 5, 10, 10, 5, 1)/32$. Recall that this indicates the proportion of time that is spent in the various states. Suppose first that the process starts in state 0. The after an even number of steps it must be in an even-numbered state, while after an odd number of steps it cannot be in an even-numbered state. So the probability is *twice* the probability in the equilibrium distribution, but for the even-numbered states only. This gives

$$16P^{2n} \to \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array} \begin{array}{c} \begin{array}{cccccc} 0 & 1 & 2 & 3 & 4 & 5 \end{array} \\ \begin{pmatrix} 1 & 0 & 10 & 0 & 5 & 0 \\ 0 & 5 & 0 & 10 & 0 & 1 \\ 1 & 0 & 10 & 0 & 5 & 0 \\ 0 & 5 & 0 & 10 & 0 & 1 \\ 1 & 0 & 10 & 0 & 5 & 0 \\ 0 & 5 & 0 & 10 & 0 & 1 \end{pmatrix} \end{array}.$$

The limit of the odd powers of P is similar except that the two kinds of row are interchanged.

Note that we can reorder the rows and columns to get the following forms:

$$16P^{2n} \to \begin{array}{c} \\ 0 \\ 2 \\ 4 \\ 1 \\ 3 \\ 5 \end{array} \begin{array}{c} \begin{array}{cccccc} 0 & 2 & 4 & 1 & 3 & 5 \end{array} \\ \begin{pmatrix} 1 & 10 & 5 & 0 & 0 & 0 \\ 1 & 10 & 5 & 0 & 0 & 0 \\ 1 & 10 & 5 & 0 & 0 & 0 \\ 0 & 0 & 0 & 5 & 10 & 1 \\ 0 & 0 & 0 & 5 & 10 & 1 \\ 0 & 0 & 0 & 5 & 10 & 1 \end{pmatrix} \end{array} \qquad 16P^{2n+1} \to \begin{array}{c} \\ 0 \\ 2 \\ 4 \\ 1 \\ 3 \\ 5 \end{array} \begin{array}{c} \begin{array}{cccccc} 0 & 2 & 4 & 1 & 3 & 5 \end{array} \\ \begin{pmatrix} 0 & 0 & 0 & 5 & 10 & 1 \\ 0 & 0 & 0 & 5 & 10 & 1 \\ 0 & 0 & 0 & 5 & 10 & 1 \\ 1 & 10 & 5 & 0 & 0 & 0 \\ 1 & 10 & 5 & 0 & 0 & 0 \\ 1 & 10 & 5 & 0 & 0 & 0 \end{pmatrix} \end{array}.$$

The rest of this chapter will be concerned with *random walks*, a special kind of Markov chain.

Example 7. Unrestricted simple random walk on the integers. There are infinitely many states indexed by the integers. As before Z_t is used to denote the state at time t. The process moves in independent steps X_t just before time t, where X_t may be ± 1 or 0 with probability p, q, r such that $p + q + r = 1$. Let Z_0 be the initial state. Then $Z_t = Z_0 + \sum_{s=1}^{t} X_s$. We assume that p, q are strictly positive, but r may be 0. The mean step size is $p-q$ and the variance of step size is $p + q - (p-q)^2 = v$, say. The total displacement after n steps has mean $n(p-q)$ with variance nv. If n is large, the Central Limit Theorem applies. Two cases arise:

(a) $p \neq q$, say $p > q$, then the mean displacement is $O(n)$ and the standard deviation is $O(\sqrt{n})$. In this case the process drifts off to infinity. The spread increases more slowly than the mean. For any a, $\mathrm{P}(Z_n \leq a) \to 0$ as $n \to \infty$.

(b) $p = q$ Here the mean displacement is always 0, but the standard deviation is $\sqrt{2pn} = O(\sqrt{n})$. Suppose the initial position is 0, then for any a, $\mathrm{P}(Z_n \le a) = \Phi(a/\sqrt{2pn}) \to 1/2$. This process makes long excursions in both directions, returning to the starting point occasionally. However it does not 'drift off to infinity' in either direction.

We turn next to questions of the type: How long does it take to return to the starting point, or to reach a certain point? We now introduce *absorbing barriers*, as in Example 3. The motion is as before until Z_t becomes equal to a or b. (Assume $a < b$.) Once $Z_t = a$ or b, all future steps are 0, i.e. the process stays for ever at the barrier and it has been *absorbed* there. Suppose the starting position is k, where $a < k < b$. We shall set up *difference equations* for the probability of being absorbed at a starting from k, which will be denoted $p_a(k)$. This is done by a *first step analysis*. Starting from k there are three possibilities for the first step, namely 1, -1, 0 with probability p, q, r. If the first step is 1 then the subsequent probability of being absorbed at a is $p_a(k+1)$. Similarly for the other possibilities. Putting the probabilities together gives the difference equation

$$p_a(k) = pp_a(k{+}1) + qp_a(k{-}1) + rp_a(k) \qquad (a < k < b) \qquad (*)$$
$$\Big(= \sum_j \mathrm{P}(\text{first step to } j)\mathrm{P}(\text{absorption at } a \mid \text{first step to } j)\Big).$$

This equation relates $p_a(k)$, $p_a(k{+}1)$ and $p_a(k{-}1)$. We shall solve for all relevant values of k. We need also some *boundary conditions*. If the starting position is a then absorption at a has already occurred, so $p_a(a) = 1$, while if the starting position is b, absorption at b has occurred so that absorption at a cannot occur. Thus $p_a(b) = 0$.

The method is similar to that used for second order linear differential equations. The general solution to (*) is found by solving the *auxiliary equation*

$$\lambda^k = p\lambda^{k+1} + q\lambda^{k-1} + r\lambda^k.$$

The roots are $\lambda = 1$, p/q. There are two distinct roots unless $p = q$.

Exercise. Check that, in the case of $p \neq q$, $A + B(q/p)^k = p_a(k)$ satisfies (*), where A, B are arbitrary constants. In the case of $p = q$ the general solution is $A + Bk$.

Finally we use the boundary conditions to determine the appropriate values of A, B. The solutions are

$$\text{in case } p \neq q, \quad \frac{(q/p)^b - (q/p)^k}{(q/p)^b - (q/p)^a}; \qquad \text{in case } p = q, \quad \frac{b-k}{b-a}.$$

You should check that these have the correct form and satisfy the boundary conditions.

Note that, for the case of $p = q$ the probabilities as a function of k lie on a straight line.

If $q > p$ the points lie on a curve that is concave, i.e. lies above any chord. Absorption at a is likely unless the starting point is very close to b. On the whole the movement tends to be to the left.

Equation (*) also applies to the probability of absorption at b, but the boundary conditions are changed. It can be shown that the solutions are

$$\text{in case } p \neq q, \quad \frac{(q/p)^k - (q/p)^a}{(q/p)^b - (q/p)^a}; \qquad \text{in case } p = q, \quad \frac{k-a}{b-a}.$$

Note that in all cases $p_a(k) + p_b(k) = 1$, i.e. there is probability 0 that the process is not eventually absorbed.

We can use these results to get the limiting behaviour of the transition matrix in Example 3. If $p = q$ it is

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ .75 & 0 & 0 & 0 & .25 \\ .5 & 0 & 0 & 0 & .5 \\ .25 & 0 & 0 & 0 & .75 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

Time until absorption. Denote the mean time till absorption starting from k by $T(k)$ for $a \leq k \leq b$. We assume that this is finite. Again we set up difference equations. The first step analysis proceeds as before, but this time we are dealing with time and we note that one time unit is used up in making the first step. The equation obtained is

$$T(k) = pT(k+1) + qT(k-1) + rT(k) + 1 \qquad (a < k < b) \qquad (**)$$
$$(= \text{expected further time after the first step } + 1).$$

The boundary conditions are $T(a) = T(b) = 0$.
The general solution to the homogeneous equation (*) (complementary function) is the same as before. We need to find a particular solution. To do this in the case of $p \neq q$, try the function ck, where c is a constant to be determined. Substitution for $T(k)$ in (**) gives $c = 1/(q-p)$. In the case of $p = q$ the function ck satisfies (*), so try ck^2 instead. This time we get the solution $c = -1/(2p)$.

You should check that for the case of $p = q$, $T(k) = \frac{(k-a)(b-k)}{2p}$.

Example 8. A person has £9 and is very keen to increase this sum to £10. There is an opportunity to place a stake and then play a game in which the probability of winning is 0.4 . If the player wins (s)he receives an amount equal to the stake in addition to having the stake repaid. In the event of losing the stake is lost. The player is allowed to choose the stake (in multiples of 10p). The games are independent. How should the player choose the stake to maximise the probability of attaining a capital of £10?

First consider this as a random walk with $p = 0.4$, $q = 0.6$. Consider placing a stake of £1 each time. Represent the state as the capital in £. The initial position is 9. We place absorbing barriers at 0 and 10 and we are required to find the probability of absorption at 10. Using the formula above, we get $(1.5^9 - 1.5^0)/(1.5^{10} - 1.5^0) = 0.661$. The expected value of the final sum is £6.61 . Now consider betting 10p each time. Represent the state as the amount of money held as a multiple of 10p. The initial state is 90 and the absorbing barriers are at 0 and 100. The probability of reaching £10 is now 0.017 and the expected final capital is only 17p.

This leads to the idea of betting as large an amount as we can, more precisely, whichever is less of the capital and the amount by which the capital is short of £10. Initially bet £1. If the first game is lost, the capital will stand at £8 so bid £2 the next time. If the game is lost, the capital will be £6, and bid £4 next. If this is lost then only £2 remain so bid £2. If the game is lost, no money remains, but if it is won the capital is £4 and £4 should be bid. This is a Markov chain with states $\{0, 2, 4, 6, 8, 9, 10\}$, and states 0 and 10 are absorbing. We can perform a first step analysis. Use notation p_k to denote the probability of reaching 10 starting from k. We get the following equations.

$$p_9 = .4 + .6p_8 \quad p_8 = .4 + .6p_6 \quad p_6 = .4 + .6p_2 \quad p_2 = .4p_4 \quad p_4 = .4p_8$$

Solving these gives $p_9 = .807$, and the expected final amount is £8.07 .

Exercise. Use first step analysis to find the mean times till absorption in Example 5.

[Answer: Starting from 0, $\frac{p^{-5}-1}{1-p}$.]

Work out the numerical values of these times from each starting point for $p = 1/2,\ 2/3$.