IEOR 6711, HMWK 3 Solutions, Professor Sigman

1. Consider a positive recurrent Markov chain with limiting stationary distribution π. We know that π_i is, by definition, the long-run proportion of time that the chain moves into state i; $\pi_i = \lim_{n\to\infty} \frac{1}{n}\sum_{k=1}^{n} I\{X_k = i\}$. It is also a rate: The long-run rate (number of times per unit time) that the chain moves into state i.

   (a) Argue that π_i is also the long-run proportion of time (rate) that the chain moves out of state i.

   SOLUTION: Every time the chain visits state i, it must leave state i (at some time soon after) so as to be able to return to it again (which it must by recurrence). Thus there is a one-to-one correspondence between visits into state i and visits out of state i. In fact the number of visits into state i during the first n units of time (denote by $N_i^+(n)$) is equal to the number of visits out of state i (denote by $N_i^-(n)$), plus or minus 1: $N_i^+(n) = N_i^-(n) \pm 1$ Dividing by n and taking the limit as $n \to \infty$ thus yields the two limits (rates) as identical. (This basic fact has nothing to do with Markov chains, it is just a basic fact about functions/paths.)

   So, “the rate into a state must equal the rate out of that state”

   (b) Argue that $\pi_i P_{i,j}$ is the long-run rate that the chain moves from state i into state j.

   SOLUTION: π is the rate out of state i, and *independent of the past*, whenever the chain is in state i it moves next to state j with probability $P_{i,j}$ by the Markov property.

   (c) Argue that

   $$\sum_i \pi_i P_{i,j}$$

   is the long-run rate that the chain moves into state j.

   SOLUTION: Summing up (b) over all states i yields, for each fixed $j \in \mathcal{S}$, that the long-run rate into state j is

   $$\sum_i \pi_i P_{i,j}.$$

   (d) Use the above to conclude that, for each j,

   $$\pi_j = \sum_i \pi_i P_{i,j};$$

   that is, $\pi = \pi P$.

   SOLUTION: It says that “The rate out of state j equals the rate into state j, for each j” which we know must be true from (a).

2. George has $r \geq 3$ umbrellas distributed between home and office as follows: When departing home at the beginning of a day, if it is raining, then he takes an umbrella (if there is one) with him to the office. Similarly, when departing the office at the end of a day, if it is raining, then he takes an umbrella (if there is one) with him to home. Assume that independent of the past there is a fixed probability $0 \leq p \leq 1$ that it rains any time he departs a location (home or office).

   (a) Argue that X_n = the number of umbrellas at the current location JUST BEFORE he departs for the n^{th} time forms a MC, with state space $\mathcal{S} = \{0, 1, 2, 3, \ldots, r\}$ and

find the transition matrix. (For example, if $r = 3$, then $P_{2,2} = p$, $P_{2,1} = 1 - p$; $P_{0,3} = 1$ and so on.)
SOLUTION: The Markov property holds due to the independence of the past assumption on the rain, and the fact that if $X_n = i$, then we know exactly how many umbrellas are at each location (i and $r - i$), and that is all we need to know to predict the future.

$$P_{0,r} = 1; \ P_{i,r-i} = 1 - p, \ P_{i,r-i+1} = p, \ i = 1, \ldots, r.$$

(b) Explain why the stationary distribution π must exist for this chain and be unique (e.g., name a Theorem , etc.).
SOLUTION: This is a finite state space irreducible MC, hence it has a unique stationary distribution by Theorem 2.2 from Lecture Notes 4

(c) Show that π is given by $\pi_0 = (1 - p)/(r + 1 - p)$, $\pi_i = (r + 1 - p)^{-1}$, $1 \le i \le r$.
SOLUTION:
We need only verify that the given solution satisfies $\pi = \pi P$ which is easily seen to be so; we do not need to derive the solution from scratch (it was already given to us).
These equations reduce to:

$$\begin{aligned}
\pi_r &= \pi_0 + \pi_1 p \\
\pi_j &= \pi_{r-j}(1 - p) + \pi_{r-j+1} p, \ j = 1, \ldots, r - 1 \\
\pi_0 &= \pi_r (1 - p),
\end{aligned}$$

and it is easily verified that they are satisfied when the given solution is plugged in. Note how, conditional on being strictly > 0, the stationary distribution is the discrete uniform over the integers $\{1, 2, \ldots, r\}$.

(d) George gets wet if and only if he departs a location when it is raining and all the umbrellas are at the other location. What is the long-run proportion of times that George gets wet ?
SOLUTION:
$p\pi_0 = p(1 - p)/(r + 1 - p)$.

(e) What value of p *maximizes* the long-run proportion of times that George gets wet? Compute for the case when $r = 3$.
SOLUTION: ($r = 3$). Using $f(p) = p(1 - p)/(4 - p)$, we need to solve $f'(p) = 0$:

$$f'(p) = \frac{(4 - p)(1 - 2p) - (-1)(p(1 - p))}{(4 - p)^2} = \frac{p^2 - 8p + 4}{(4 - p)^2} = 0;$$

solving the quadratic $p^2 - 8p + 4 = 0$ yields $p = 4 - 2\sqrt{3} \approx 0.536$ (the other root is > 1, too large to be a probability, so we discard it). We know it is a maximum because the minimum is achieved when $p = 1$ or $p = 0$ (the boundary points).

3. *Martingale MC:* Consider the MC with state space $\mathcal{S} = \{0, 1, 2, 3, 4\}$ and transition matrix

$$P = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 1/2 & 0 & 1/2 & 0 & 0 \\ 0 & 3/5 & 0 & 1/5 & 1/5 \\ 1/10 & 1/10 & 1/10 & 1/10 & 6/10 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

(a) Show that $E(X_{n+1}|X_n = i) = i,\ i \in \mathcal{S}$. (e.g., each row i of P has mean i.) Thus, being so for each i, we conclude that $E(X_{n+1}|X_n) = X_n$: Given the present state X_n, the expected value of the future is X_n itself.

SOLUTION: The key point here is that this matrix has a very special property: Each row i forms a probability distribution on $\mathcal{S}$ that has expected value i: For each $i \in \{0, 1, 2, 3, 4\}$

$$E(X_{n+1}|X_n = i) = \sum_{j=0}^{4} jP_{i,j} \ = \ i.$$

(b) Now use the additional fact that $\{X_n\}$ is a MC to deduce that

$$E(X_{n+1}|X_n, X_{n-1}, \ldots, X_0) = X_n,\ n \geq 0:$$

Given the present state, X_n, the expected value of the future is X_n itself, *independent of the past.*

SOLUTION: Being a MC, the future is independent of the past given the present state, in particular,
$E(X_{n+1}|X_n, X_{n-1}, \ldots, X_0) = E(X_{n+1}|X_n)$. Since $E(X_{n+1}|X_n) = X_n$ from (a), the result follows.

4. *Infinite dimensional Markov chain.* Let $\{X_n : n \geq 0\}$ be any stochastic process, and let $\mathbf{Y}_0 = \{X_n : n \geq 0\}$, $\mathbf{Y}_1 = \{X_{1+n} : n \geq 0\} = \{X_1, X_2, \ldots\}, \ldots,$ $\mathbf{Y}_k = \{X_{k+n} : n \geq 0\} = \{X_k, X_{k+1}, \ldots\}$. Each $\mathbf{Y}_k$ is thus the entire infinite sequence from time k onwards. Argue that in fact $\{\mathbf{Y}_k : k \geq 0\}$ satisfies the Markov property: The future $\mathbf{Y}_{k+1}$ is independent of the past $\mathbf{Y}_0, \mathbf{Y}_1, \ldots \mathbf{Y}_{k-1}$ given the present state $\mathbf{Y}_k$. Thus, in principle, any stochastic process can be modeled as a Markov chain by including enough information (the infinite future). (But, the complexity is so high that this is of little practical value.)

SOLUTION:

If we condition on $\mathbf{Y}_k = \{X_k, X_{k+1}, \ldots\}$, then it already contains (completely determines) $\mathbf{Y}_{k+1} = \{X_{k+1}, X_{k+2}, \ldots\}$. Define the "shift operator" θ via $\theta\mathbf{Y}_k = \mathbf{Y}_{k+1}$.

$$\begin{aligned} P(\mathbf{Y}_{k+1} \in A\ |\mathbf{Y}_k, \mathbf{Y}_{k-1}, \ldots, \mathbf{Y}_0) &= P(\theta\mathbf{Y}_k \in A\ |\mathbf{Y}_k, \mathbf{Y}_{k-1}, \ldots, \mathbf{Y}_0) \\ &= I\{\theta\mathbf{Y}_k \in A\} \\ &= P(\mathbf{Y}_{k+1} \in A \mid \mathbf{Y}_k). \end{aligned}$$

Each conditional probability is just an indicator function, 1 or 0.

5. Consider a chess board (64 squares). A knight is initially in the upper left corner and is the only piece on the board (recall that a knight can only move in "L" shaped moves). Assume that every unit of time the knight moves by randomly choosing from its available legal moves in an "equally likely" fashion (independent of any past moves). What is the expected amount of time (moves) until the knight returns back to the upper left corner?

SOLUTION: Letting X_n denote the position of the knight after his n^{th} move yields an irreducible finite state Markov chain on $\mathcal{S} = \{1, 2, 3, \ldots, 64\}$. Thus it must be positive recurrent with stationary distribution π satisfying $\pi = \pi P$ (by say Theorem 2.2 from Lecture Notes 4). We want $E(\tau_{j,j}) = \pi_j^{-1}$ for $j = 1$ denoting the upper left corner state. Note that in the game of chess, a knight can move either two squares horizontally and one vertically or two vertically followed by one square horizontally, as long as it remains on the board. There are thus eight possible moves that the knight can make:

```
2 left, 1 down
2 left, 1 up
2 right, 1 down
2 right, 1 up
2 up, 1 left
2 up, 1 right
2 down, 1 left
2 down, 1 right
```

The Figure 1 below gives the number of possible moves from each position.

2	3	4	4	4	4	3	2
3	4	6	6	6	6	4	3
4	6	8	8	8	8	6	4
4	6	8	8	8	8	6	4
4	6	8	8	8	8	6	4
4	6	8	8	8	8	6	4
3	4	6	6	6	6	4	3
2	3	4	4	4	4	3	2

Figure 1: **Chessboard**

This Markov chain can be represented as a random walk on a weighted graph which has 64 nodes (we label the nodes row-by-row from left to right, $1-64$). If the knight can move from position i to j in one move, then we place an arc between i and j and define $w_{i,j}=1$; if the knight cannot movie from i to j in one move, then we do not place an arc between the two and we set $w_{i,j}=0$. The positive weights are 1 because the knight moves from node i equally likely to all admissible nodes j by assumption. From the theory of time- reversible Markov chains, we know that $\{X_n\}$ is time reversible and the stationary distribution π can be computed as follows:

$$\pi_i = \frac{\sum_{k=1}^{64} w_{i,k}}{\sum_{j=1}^{64}\sum_{k=1}^{64} w_{j,k}}, \quad i=1,\ldots,64.$$

Note that since all admissible arcs have weight 1, we need simply count the number of possible moves from each position i.

Therefore, if we label the top-left corner as position 1, then we have

$$\pi_1 = \frac{2}{2 \times 4 + 3 \times 8 + 4 \times 20 + 6 \times 16 + 8 \times 16} = \frac{1}{168}.$$

If we let $\tau_{1,1}$ be the the first passage time to 1 starting with 1, i.e., $\tau_{1,1} \equiv \inf\{n \geq 1 : X_n = 1\}$ with $X_0 = 1$, then $E[\tau_{1,1}] = 1/\pi_1 = 168$.

6. Consider the Gambler's ruin Markov chain $\{X_n\}$ on $\mathcal{S} = \{0, 1, \ldots, N\}$. $P_{0,0} = 1 = P_{N,N}$, while otherwise $P_{i,1+1} = p$, $P_{i,i-1} = 1-p$, $1 < i < N$. Suppose that $X_0 = i$, for $1 < i < N$. Let $X(i) = \lim_{n\to\infty} X_n$. Explain why the rv $X(i)$ must exist (with probability 1), and give its probability mass function: $p_k = P(X(i) = k),\ k \in \{0, 1, \ldots, N\}$.

   SOLUTION:

   From the gambler's ruin problem we know that if $X_0 = i$, then the chain will eventually either hit N with probability $P_i(N)$ or hit 0 with probability $1 - P_i(N)$. Thus $P(X(i) = 0) = 1 - P_i(N),\ P(X(i) = N) = P_i(N),$ and $P(X(i) = k) = 0,\ 1 < k < N$.