IEOR 6711, HMWK 2, Professor Sigman

1. For the gambler's ruin problem with $p = 0.60$ and $N = 3$: If the gambler starts off with \$1, then what is the probability that the game ends in no more than 3 gambles? (The game ends when either the gambler reaches N or goes broke, whichever happens first. So we want the probability that the simple random walk MC chain associated with the gambler's ruin problem ($P_{00} = P_{N,N} = 1$), having started with $X_0 = 1$, is either in state 0 or state N at time $n = 3$.)

   SOLUTION:

   Compute P^3 for the transition matrix, and then sum up $P^{(3)}_{1,0} + P^{(3)}_{1,N}$ for the desired answer.

2. (*Easy proof of Gambler's ruin problem in the symmetric case*) Consider the simple *symmetric* ($p = 0.5$) random walk, $R_n = \Delta_1 + \cdots + \Delta_n$, with $R_0 = i$ where $0 < i < N$. Let $\tau = \min\{n \geq 1 : R_n \in \{0, N\}\}$, the time at which the gambler stops gambling in the gambler's ruin problem.

   (a) Explain why $E(\tau) < \infty$ using elementary MC ideas.

   SOLUTION:

   One quick way based on what you already proved on HMWK 1: Change the random walk to the one associated with the gambler's ruin problem, that is, in which there are the two absorbing (recurrent) states 0 and N; $P_{0,0} = P_{N,N} = 1$. The other states $\{1, 2, \ldots, N-1\}$ are all transient. Recall the square matrix $S = (s_{ij})$ for transient states, where $s_{i,j}$ = the expected total number of visits to transient state j if X_0 =transient state i, and is finite for each such pair.

   ($S = (I - P_T)^{-1}$.) Summing all of them (over j) for any fixed i is thus also finite and is in fact $= E(\tau)$; if $X_0 = i$, then

   $$E(\tau) = \sum_{j=1}^{N-1} s_{i,j} < \infty$$

   (b) Derive the identity

   $$E\Big[i + \sum_{n=1}^{\tau} \Delta_n\Big] = NP_i(N).$$

   SOLUTION:

   Since $R_0 = i$, we have $R_\tau = i + \sum_{n=1}^{\tau} \Delta_n$. But by definition, R_τ is either N with probability $P_i(N)$ or 0 with probability $1 - P_i(N)$; hence leading to expected value $E(R_\tau) = NP_i(N)$.

   (c) Use the above identity (justifying along the way) to show that $P_i(N) = i/N$.

   SOLUTION:

   We apply Wald's equation;

   $$E \sum_{n=1}^{\tau} \Delta_n = E(\tau)E(\Delta) = 0,$$

   using the fact that $E(\Delta_n) = 0$: thus $i = NP_i(N)$, and we solve $P_i(N) = i/N$.

3. Consider a man stuck in a cave that has 3 doors: he is equally likely to choose any such door (1/3 probability) with the objective being to reach freedom (escape). If he chooses door 1, he wanders for 5 days and returns to the cave. If he chooses door 2, he wanders for 3 days and returns to the cave. If he chooses door 3, he travels for 2 days and reaches freedom. Whenever choosing the wrong door (1,2) hence returning to the cave, he is (independent of the past) again equally likely to choose any such door.

   (a) Let X denote the total number of days until he reaches freedom. Let D denote the first door he chooses. By conditioning on $D = i,\ i = 1, 2, 3$, derive the following equation:
   $E(X) = 1/3((5 + E(X)) + (3 + E(X)) + 2)$.
   SOLUTION: From the independent of the past assumption: $E(X \mid D = 1) = 5 + E(X),\ E(X \mid D = 2) = 3 + E(X),\ E(X \mid D = 3) = 2$. Now use

   $$E(X) = \sum_{i=1}^{3} E(X \mid D = i)P(D = i),$$

   where $P(D = i) = 1/3,\ i = 1, 2, 3$.

   (b) Solve for $E(X)$.
   SOLUTION: $E(X) = 10$.

   (c) Here you are to use Wald's equation to find $E(X)$: Define an iid sequence of rvs $Y_1, Y_2, \cdots$ and a stopping time τ, such that

   $$X = \sum_{n=1}^{\tau} Y_n.$$

   Now use this to compute $E(X)$.
   SOLUTION:
   Y_i = the number of days journey after choosing a door for the i^{th} time. $P(Y = 5) = P(Y = 3) = P(Y = 2) = 1/3$, and so $E(Y) = 10/3$. $\tau = \min\{n \geq 1 : Y_n = 3\}$, hence geometric with mean 3. $E(X) = E(\tau)E(Y) = 10$.

4. Suppose that Y_1 and Y_2 are *identically distributed* rvs on the non-negative integers, and their joint distribution is $P(Y_1 = k, Y_2 = l) = a_{k,l},\ k, l \geq 0$. Show how to construct a stationary sequence $\{X_n : n \geq 0\}$ of rvs, such that the joint distribution (X_n, X_{n+1}) is the same as that of (Y_1, Y_2) for all $n \geq 0$. (By stationary, we mean in particular that X_n has the same distribution for all $n \geq 0$.)

   SOLUTION: Define a Markov chain with transition probabilities given by
   $P_{i,j} = P(Y_2 = j|Y_1 = i),\ i, j \geq 0$, and start it off initially with distribution $P(X_0 = j) = P(Y = j),\ j \geq 0$, where Y is the common distribution of Y_1, Y_2:

   $$P(Y = j) = P(Y_1 = j) = P(Y_2 = j) = \sum_{i=0}^{\infty} a_{j,i}.$$

   By construction, this is a Markov chain with the desired properties: It is stationary because X_n has the same distribution as X_0 for all $n \geq 0$ (namely the distribution of Y), and thus (X_n, X_{n+1}) has the same distribution as (Y_1, Y_2) due to the definition of the transition probabilities.

5. Recall the geometric series

$$(1-x)^{-1} = \sum_{n=0}^{\infty} x^n,$$

for $|x| < 1$.

Let $M = (m_{ij})$ be a $b \times b$ matrix and suppose that $M^n \to 0$ as $n \to \infty$ (where here 0 denotes the 0 matrix, and the convergence means that $m_{ij}^{(n)} \to 0$ for each fixed pair i, j). Then

$$(1-M)^{-1} = \sum_{n=0}^{\infty} M^n. \tag{1}$$

This is so because

$$(I + M + M^2 + M^3 + \cdots + M^n)(I - M) = I - M^{n+1},$$

and if $M^n \to 0$ then as $n \to \infty$, the right hand side converges to I yielding the result.

Let $M = P_T$, the transient states matrix from Homework 1, Problem 4. Argue that $P_T^n \to 0$. Hint: use Chapman-Kolmogorov to see that $P_T^{(n)} = P_T^n$ just as $P^{(n)} = P^n$, and observe that $m_{ij}^{(n)} = p_{i,j}^{(n)} = P(X_n = j \mid X_0 = i)$.

(1) is useful since we can thus approximate $(I - P_T)^{-1}$ by choosing a large k and using

$$(I - P_T)^{-1} \approx \sum_{n=0}^{k} M^n.$$

SOLUTION: As is easily checked, indeed $P_T^{(n)} = P_T^n$, and since all the states $i, j \in T$ are transient, it must hold that if $X_0 = i \in T$, then $X_n \notin T$ for sufficiently large n wp1. Thus $p_{i,j}^{(n)} = P(X_n = j \mid X_0 = i) \to 0$ as $n \to \infty$; that is, $P_T^n \to 0$.

6. Consider an irreducible Markov chain $\{X_n\}$ with finite state space $\mathcal{S} = \{1, 2, \ldots N\}$, and transition matrix P. Consider the set of N equations, $\pi = \pi P$, and the additional equation $\pi_1 + \cdots + \pi_N = 1$. Thus a total of $N + 1$ equations. Show that if you remove any one of the first N equations, then in fact, the remaining N equations still uniquely determine the stationary distribution π (e.g., they have a unique probability solution that IS the stationary distribution π).

SOLUTION:

Without loss of generality, let us suppose that it is the first equation $\pi_1 = \sum_{i=1}^{N} \pi_i P_{i,1}$ that we wish to replace with $\sum_{j=1}^{N} \pi_j = 1$. It then suffices to show that the remaining $N - 1$ equations can be combined to reproduce equation 1.

To this end, we simply sum them up yielding

$$\begin{aligned} \pi_2 + \cdots + \pi_N &= \pi_1 (\sum_{j=2}^{N} P_{1,j}) + \cdots + \pi_N (\sum_{j=2}^{N} P_{N,j}) \\ &= \pi_1 (1 - P_{1,1}) + \cdots + \pi_N (1 - P_{N,1}). \end{aligned}$$

Now using $\sum_{j=1}^{N} \pi_j = 1$ on both sides yields

$$1 - \pi_1 = 1 - \sum_{i=1}^{N} \pi_i P_{i,1},$$

or

$$\pi_1 = \sum_{i=1}^{N} \pi_i P_{i,1},$$

as was to be shown.

Since the chain is assumed irreducible, and the state space is finite, there is a unique probability solution to $\pi = \pi P$; the proof is complete.

7. A Markov chain transition matrix P is called *doubly stochastic* if its *columns* each sum to 1. Thus each row and each column are non-negative elements summing to 1. Show that if a Markov chain has a finite state space $\mathcal{S} = \{1, 2, \ldots, b\}$ and is both irreducible and doubly stochastic, then the stationary distribution π is the discrete uniform distribution over the set $\mathcal{S}$; $\pi_j = 1/b, \; j \in \{1, 2, \ldots, b\}$.

   SOLUTION:

   Since the chain is irreducible, it suffices to simply "plug in" the proposed solution for π to verify that it satisfies $\pi = \pi P$; since it does so, it is the unique stationary distribution: For each $j, \in \{1, 2, \ldots, b\}$,

$$1/b = \pi_j = \sum_{i=1}^{b} (1/b) P_{j,1} = 1/b,$$

   since $\sum_{i=1}^{b} P_{i,j} = 1$.

8. Let $Y_n = \sum_{k=1}^{n} D_k, \; n \geq 1$, where the $\{D_k\}$ are iid with a discrete uniform distribution over the set $\{1, 2, 3, 4, 5, 6\}$ (e.g., the outcome of rolling a dice).

   Find

$$\lim_{N \to \infty} \frac{1}{N} \sum_{n=1}^{n} P(Y_n = \text{a multiple of } 13).$$

   (HINT: Use the above problem concerning doubly stochastic.)

   SOLUTION: Let X_n denote the value of Y_n modulo 13. That is, X_n is the remainder when Y_n is divided by 13. X_n is an irreducible Markov chain with states $0, 1, ...12$, and note that $X_n = 0$ if and only if Y_n is a multiple of 13. Thus we want to compute π_0 for this Markov chain. It is easy to verify that $\Sigma_i P_{ij} = 1$ for all j, that is, the chain is doubly stochastic. (For instance, for $j = 3$, $\Sigma_i P_{ij} = P_{2,3} + P_{1,3} + P_{0,3} + P_{12,3} + P_{11,3} + P_{10,3} = \frac{1}{6} + \frac{1}{6} + \frac{1}{6} + \frac{1}{6} + \frac{1}{6} + \frac{1}{6} = 1$.) Hence from the previous problem, we know that $\pi_0 = \frac{1}{13}$.