Stoch. Proc. HW assignment 12. Due Tuesday, December 12

1. Let us consider the discrete-time birth/death process (X_n) with state space $\mathbb{N}$ and transition rules

$$\begin{aligned}\mathbb{P}(\, X_{n+1} = x+1 \,|\, X_n = x\,) &= p_x, \quad x = 1,2,3,\dots\\ \mathbb{P}(\, X_{n+1} = x-1 \,|\, X_n = x\,) &= q_x, \quad x = 1,2,3,\dots,\\ \mathbb{P}(\, X_{n+1} = x \,|\, X_n = x\,) &= 1 - p_x - q_x, \quad x = 1,2,3,\dots\\ \mathbb{P}(\, X_{n+1} = 0 \,|\, X_n = 0\,) &= 1.\end{aligned}$$

Let us define the function $\alpha : \mathbb{N} \to \mathbb{R}$ by

$$\alpha(x) = \sum_{y=0}^{x-1} \frac{q_1 q_2 \dots q_y}{p_1 p_2 \dots p_y}. \qquad \text{(Note: } \alpha(0) = 0 \text{ and } \alpha(1) = 1\text{)}$$

Show that (M_n) is a martingale, where $M_n := \alpha(X_n)$. *Hint:* See page 177 of scanned lecture notes.

Solution: We only need to show that $\alpha : \mathbb{N} \to \mathbb{R}$ is a harmonic function, i.e.,

$$\alpha(x) \equiv q_x \alpha(x-1) + (1 - p_x - q_x)\alpha(x) + p_x \alpha(x+1).$$

Rearranging this, we only need to show that

$$q_x \cdot (\alpha(x) - \alpha(x-1)) \equiv p_x \cdot (\alpha(x+1) - \alpha(x))\,.$$

This is indeed true, because

$$q_x \cdot (\alpha(x) - \alpha(x-1)) = q_x \cdot \frac{q_1 q_2 \dots q_{x-1}}{p_1 p_2 \dots p_{x-1}} = p_x \cdot \frac{q_1 q_2 \dots q_x}{p_1 p_2 \dots p_x} = p_x \cdot (\alpha(x+1) - \alpha(x))\,.$$

Note that the $x = 0$ case is a bit different: in this case the equation $\alpha(0) = \sum_{y \in \mathbb{N}} p(0,y)\alpha(y)$ reduces to $\alpha(0) = \alpha(0)$, which trivially holds.

Remark: We have already solved this exercise on page 45-46 of the scanned lecture notes.

2. You are in a casino. If you bet 1 dollar in the n'th round then your your net profit in that round is ξ_n dollars where $\mathbb{P}(\xi_n = +1) = p$, $\mathbb{P}(\xi_n = -1) = q$, $q+p = 1$, $p > 1/2$ and $\xi_1, \xi_2, \ldots$ are i.i.d. In other words: with probability $q < 1/2$ you lose the bet and with probability $p = 1-q > 1/2$ you double your bet. You bet C_n dollars in the n'th round. Obviously, your betting strategy has to be *predictable*, i.e., when you make your n'th bet, you do not yet know the value of ξ_n, so let us assume that C_n is measurable w.r.t. $(\xi_1, \ldots, \xi_{n-1})$. Denote by y_0 your initial wealth and by Y_n your total wealth after the end of the n'th round. You cannot go in debt, so let's assume $0 \le C_n \le Y_{n-1}$, $n > 0$. You are allowed to play N rounds.

   Your goal is to maximize your expected rate of return $\mathbb{E}(\log Y_N - \log y_0)$.

   (a) Define
   $$f(x) := p \ln(1+x) + (1-p)\ln(1-x), \qquad 0 \le x \le 1.$$
   Show that f is strictly concave. Find $\max_{0 \le x \le 1} f(x)$ and show that $\max_{0 \le x \le 1} f(x) > 0$.

   (b) Show that for any predictable betting strategy we have $\mathbb{E}\big(\ln(Y_{n+1}) \,|\, (\xi_1, \ldots, \xi_n)\big) = \ln(Y_n) + f(\frac{C_{n+1}}{Y_n})$.

   (c) Let $Z_n := \log Y_n - n\alpha$, where $\alpha = p \log p + q \log q + \log 2$. Prove that for any predictable betting strategy the process (Z_n) is a supermartingale. Show that this implies $\mathbb{E}(\log Y_N - \log y_0) \le N\alpha$.
   Hint: See page 182 of the scanned lecture notes for the definition of a *supermartingale*.

   (d) Describe the betting strategy for which Z_n is a martingale and that $\mathbb{E}(\log Y_N - \log y_0) = N\alpha$ is achieved by this strategy. (Sometimes they call this the *log-optimal portfolio* in economics)

Solution:

(a) We have $f'(x) = \frac{p}{1+x} - \frac{1-p}{1-x}$ and $f''(x) = \frac{-p}{(1+x)^2} - \frac{1-p}{(1-x)^2}$, thus f is strictly concave and attains its unique maximum at x^* where $f'(x^*) = 0$, that is $x^* = 2p - 1$. Note that $0 < x^* < 1$ follows from our assumptions.

$$\max_{0 \le x \le 1} f(x) = f(x^*) = p \ln(1 + 2p - 1) + (1-p)\ln(1 - 2p + 1) =$$
$$p \cdot (\ln(p) + \ln(2)) + q \cdot (\ln(q) + \ln(2)) = \alpha.$$

It remains to show that $\max_{0 \le x \le 1} f(x) = \alpha > 0$. This follows from the fact that the strictly convex function $p \mapsto p \log p + (1-p)\log(1-p) + \log 2$ attains its unique minimum at $p = \frac{1}{2}$, and that this minimum is equal to 0.

(b) Let us introduce $x_n = \frac{C_n}{Y_{n-1}}$, i.e., the fraction of your money that you bet in the n'th round. Since $0 \le C_n \le Y_{n-1}$, we have $0 \le x_n \le 1$. Also note:

$$Y_{n+1} = Y_n + C_{n+1}\xi_{n+1} = Y_n \cdot (1 + x_{n+1}\xi_{n+1}). \tag{1}$$

Now we calculate

$$\mathbb{E}(\ln(Y_{n+1}) \,|\, (\xi_1, \ldots, \xi_n)) \overset{(1)}{=} \mathbb{E}(\ln(Y_n) + \ln(1 + x_{n+1}\xi_{n+1}) \,|\, (\xi_1, \ldots, \xi_n)) \overset{(*)}{=}$$
$$\ln(Y_n) + p\ln(1 + x_{n+1}) + (1-p)\ln(1 - x_{n+1}) = \ln(Y_n) + f(x_{n+1}),$$

where in the equation marked by $(*)$ we used that Y_n and x_{n+1} are measurable w.r.t. $(\xi_1, \ldots, \xi_n)$ and that ξ_{n+1} is independent from $(\xi_1, \ldots, \xi_n)$.

(c) First we show that (Z_n) is a supermartingale:

$$\mathbb{E}(Z_{n+1} \,|\, (\xi_1, \ldots, \xi_n)) = \mathbb{E}(\ln(Y_{n+1}) \,|\, (\xi_1, \ldots, \xi_n)) - (n+1)\alpha =$$
$$\ln(Y_n) + f(x_{n+1}) - (n+1)\alpha \le \ln(Y_n) + \max_{0 \le x \le 1} f(x) - (n+1)\alpha = \ln(Y_n) - n\alpha = Z_n. \tag{2}$$

Since (Z_n) is a supermartingale, we have $\mathbb{E}(Z_N) \le \mathbb{E}(Z_0)$ (see page 182 of the scanned lecture notes), thus $\mathbb{E}(\ln(Y_N) - N\alpha) \le \ln(y_0)$, thus $\mathbb{E}(\log Y_N - \log y_0) \le N\alpha$.

(d) The best strategy is to bet $C_{n+1} = (2p-1)Y_n$ in round $n+1$ (which makes $x_{n+1} = x^*$), so that $f(x_{n+1}) = \max_{0 \le x \le 1} f(x) = \alpha$ and thus we have equality everywhere in (2), therefore we have $\mathbb{E}(Z_{n+1} \,|\, (\xi_1, \ldots, \xi_n)) = Z_n$, therefore (Z_n) is a martingale. This implies $\mathbb{E}(Z_0) = \mathbb{E}(Z_N)$ (see page 176 of the scanned lecture notes) and $\mathbb{E}(\log Y_N - \log Y_0) = N\alpha$.

3. There are two amoebae at time 0 in a Petri dish: one of them is red, the other one is blue. An amoeba splits after an EXP(1) distributed waiting time into two amoebae which are identical to the original one. Denote by X_n the number of blue amoebae after the n'th split. Note that $X_0 = 1$ and that the total number of amoebae in the Petri dish after the n'th split is $n+2$.

   (a) Calculate $\mathbb{P}(X_{n+1} = x \mid X_n = x)$ and $\mathbb{P}(X_{n+1} = x+1 \mid X_n = x)$ for $x \in \{1, \dots, n+1\}$.

   (b) Let us define $M_n = \frac{X_n}{n+2}$ (i.e., M_n is the fraction of blue amoebae). Show that (M_n) is a martingale.

   (c) Calculate $\mathbb{P}(X_2 = x)$ for $x = 1, 2, 3$. Calculate $\mathbb{P}(X_3 = x)$ for $x = 1, 2, 3, 4$.

   (d) Based on the pattern that you see in the previous sub-exercise, make a conjecture about the distribution of X_n for general n and verify your conjecture using induction on n.

Solution: The Markov chain (X_n) is famous, it is called *Pólya's urn model.*

(a) If $X_n = x$, then there are $n+2$ amoebae in the Petri dish and x of them are blue. Each of these $n+2$ amoebae have their own independent EXP(1) clock. Since the ringing rates are the same for all of the amoebae, each of them has the same chance of being the first one whose clock rings. Thus the one who splits next is a uniformly chosen amoeba, therefore it is blue with probability $\frac{x}{n+2}$. This gives

$$\mathbb{P}(X_{n+1} = x \mid X_n = x) = 1 - \frac{x}{n+2}, \qquad \mathbb{P}(X_{n+1} = x+1 \mid X_n = x) = \frac{x}{n+2}. \tag{3}$$

(b) We will show that (M_n) is a martingale w.r.t. the filtration of (X_n), i.e., that $\mathbb{E}(M_{n+1} \mid \underline{X}_n) = M_n$. Let η_n be the indicator of the event that n'th split produced a new blue amoeba. In other words, let $\eta_{n+1} = X_{n+1} - X_n$. It follows from (3) that

$$\mathbb{E}(\eta_{n+1} \mid \underline{X}_n) = M_n. \tag{4}$$

Now

$$M_{n+1} = \frac{X_{n+1}}{n+3} = \frac{X_n + \eta_n}{n+3} = \frac{(n+2)M_n + \eta_{n+1}}{n+3} = (1 - \frac{1}{n+3})M_n + \frac{1}{n+3}\eta_{n+1}. \tag{5}$$

In words: M_{n+1} is the weighted average of M_n and η_{n+1}. We can now conclude

$$\mathbb{E}(M_{n+1} \mid \underline{X}_n) \overset{(5)}{=} (1 - \frac{1}{n+3})\mathbb{E}(M_n \mid \underline{X}_n) + \frac{1}{n+3}\mathbb{E}(\eta_{n+1} \mid \underline{X}_n) \overset{(4)}{=} (1 - \frac{1}{n+3})M_n + \frac{1}{n+3}M_n = M_n.$$

(c) $\mathbb{P}(X_1 = 1) = \mathbb{P}(X_1 = 2) = \frac{1}{2}$.
$\mathbb{P}(X_2 = 1) = \mathbb{P}(X_1 = 1)\frac{2}{3} = \frac{1}{3}$,
$\mathbb{P}(X_2 = 2) = \mathbb{P}(X_1 = 1)\frac{1}{3} + \mathbb{P}(X_1 = 2)\frac{1}{3} = \frac{1}{3}$,
$\mathbb{P}(X_2 = 3) = \mathbb{P}(X_1 = 2)\frac{2}{3} = \frac{1}{3}$.
One can also similarly show $\mathbb{P}(X_3 = 1) = \mathbb{P}(X_3 = 2) = \mathbb{P}(X_3 = 3) = \mathbb{P}(X_3 = 4) = \frac{1}{4}$.

(d) We will prove by induction on n the statement

$$\mathbf{P}(X_n = k) = \frac{1}{n+1} \quad \text{for any} \quad n \geq 0 \quad \text{and} \quad 1 \leq k \leq n+1. \tag{6}$$

In words: we will show that X_n is uniformly distributed over the set $\{1, 2, \dots, n+1\}$.
The statement (6) is true for $n = 0$, since $\mathbf{P}(X_0 = 1) = 1$. Assuming that (6) holds for some $n \geq 0$, we will show that (6) holds for $n+1$ and any $1 \leq k \leq n+2$, i.e., we will deduce

$$\mathbf{P}(X_{n+1} = k) = \frac{1}{n+2} \quad \text{for any} \quad 1 \leq k \leq n+2. \tag{7}$$

First we show (7) for $k = 1$:

$$\mathbf{P}(X_{n+1} = 1) \overset{(3)}{=} \mathbb{P}(X_n = 1)\frac{n+1}{n+2} \overset{(6)}{=} \frac{1}{n+2}.$$

Now we show (7) for $k = n+2$:

$$\mathbf{P}(X_{n+1} = n+2) \overset{(3)}{=} \mathbb{P}(X_n = n+1)\frac{n+1}{n+2} \overset{(6)}{=} \frac{1}{n+2}.$$

Now we show (7) for $1 < k < n+2$:

$$\mathbf{P}(X_{n+1}=k) \overset{(3)}{=} \mathbf{P}(X_n = k-1)\frac{k-1}{n+2} + \mathbf{P}(X_n = k)\frac{n+2-k}{n+2} \overset{(6)}{=}$$
$$\frac{1}{n+1}\frac{k-1}{n+2} + \frac{1}{n+1}\frac{n+2-k}{n+2} = \frac{1}{n+1}\frac{n+1}{n+2} = \frac{1}{n+2}$$

We have checked that (7) holds, so by induction (6) holds.

Alternative Solution: Let us prove (6) without induction, using combinatorics.
As the new amoebae are born, let us write down the sequence of their colours. For example $BRRBR$ means that the first newborn was blue, the next one red, then red, then blue, then red. We can arrive at the event $\{X_n = k\}$ using $\binom{n}{k-1}$ different colour sequences of length n with $k-1$ blues and $n-k+1$ reds in it. We claim that every such color sequence occurs with probability $\frac{(k-1)!(n-k+1)!}{(n+1)!}$.
We first derive (6) from this claim:

$$\mathbb{P}(X_n = k) = \binom{n}{k-1}\frac{(k-1)!(n-k+1)!}{(n+1)!} = \frac{n!}{(k-1)!(n-k+1)!}\frac{(k-1)!(n-k+1)!}{(n+1)!} = \frac{1}{n+1}.$$

Now we prove the claim. We can write the probability of the occurrence of a particular colour sequence of length n as the product of n fractions, where each term of the product is the conditional probability that the next colour appears in the sequence, given the present proportion of reds and blues determined by the colours of previous births. For example, the product corresponding to $BRRBR$ is $\frac{1}{2}\frac{1}{3}\frac{2}{4}\frac{2}{5}\frac{3}{6}$. In a colour sequence of length n with $k-1$ blues and $n-k+1$ reds in it, the product of the denominators is $(n+1)!$, the product of the blue numerators is $(k-1)!$ and the product of the red numerators is $(n-k+1)!$.

Remark: More can be said about Pólya's urn model, but unfortunately our Stochastic Processes course has come to an end. If you are interested, Google *Pólya's urn model*.