Course:	Introduction to Stochastic Processes
Term:	Fall 2019
Instructor:	Gordan Žitković

Lecture 5
Random walks - advanced methods

5.1 Stopping times

Our last application of generating functions dealt with sums evaluated between 0 and some random time N. An especially interesting case occurs when the value of N depends directly on the evolution of the underlying stochastic process. Even more important is the case where time's arrow is taken into account. If you think of N as the time you *stop* adding new terms to the sum, it is usually the case that you are not allowed (able) to see the values of the terms you would get if you continued adding. Think of an investor in the stock market. Her decision to stop and sell her stocks can depend only on the information available up to the moment of the decision. Otherwise, she would sell at the absolute maximum and buy at the absolute minimum, making tons of money in the process. Of course, this is not possible unless you are clairvoyant, so the mere mortals have to restrict their choices to so-called **stopping times**.

Definition 5.1.1. Let $\{X_n\}_{n\in\mathbb{N}_0}$ be a stochastic process. A random variable T taking values in $\mathbb{N}_0 \cup \{+\infty\}$ is said to be a **stopping time** with respect to $\{X_n\}_{n\in\mathbb{N}_0}$ if for each $n \in \mathbb{N}_0$ there exists a function $G^n : \mathbb{R}^{n+1} \to \{0,1\}$ such that

$$\mathbf{1}_{\{T=n\}} = G^n(X_0, X_1, \dots, X_n), \text{ for all } n \in \mathbb{N}_0.$$

The functions G^n are called the **decision functions**, and should be thought of as a black box which takes the values of the process $\{X_n\}_{n\in\mathbb{N}_0}$ observed up to the present point and outputs either 0 or 1. The value 0 means *keep going* and 1 means *stop*. The whole point is that the decision has to based only on the available observations and not on the future ones.

Example 5.1.2.

1. The simplest examples of stopping times are (non-random) *deterministic times*. Just set $T = 5$ (or $T = 723$ or $T = n_0$ for any $n_0 \in \mathbb{N}_0 \cup \{+\infty\}$), no matter what the state of the world $\omega \in \Omega$ is. The family of decision rules is easy to construct:

$$G^n(x_0, x_1, \dots, x_n) = \begin{cases} 1, & n = n_0, \\ 0, & n \neq n_0. \end{cases}.$$

   Decision functions G^n do not depend on the values of $X_0, X_1, \dots, X_n$ *at all*. A gambler who stops gambling after 20 games, no matter of what the winnings or losses are uses such a rule.

2. Probably the most well-known examples of stopping times are *(first) hitting times*. They can be defined for general stochastic processes, but we will stick to simple random walks for the purposes of this example. So, let $X_n = \sum_{k=0}^{n} \xi_k$ be a simple random walk, and let T_l be the first time X hits the level $l \in \mathbb{N}$e More precisely, we use the following slightly non-intuitive but mathematically correct definition

$$T_l = \min\{n \in \mathbb{N}_0 \, : \, X_n = l\}.$$

   The set $\{n \in \mathbb{N}_0 \, : \, X_n = l\}$ is the collection of all time-points at which X visits the level l. The earliest one - the minimum of that set - is the first hitting time of l. In states of the world $\omega \in \Omega$ in which the level l just never gets reached, i.e., when $\{n \in \mathbb{N}_0 \, : \, X_n = l\}$ is an empty set, we set $T_l(\omega) = +\infty$. In order to show that T_l is indeed a stopping time, we need to construct the decision functions G^n, $n \in \mathbb{N}_0$. Let us start with $n = 0$. We would have $T_l = 0$ in the (impossible) case $X_0 = l$, so we always have $G^0(X_0) = 0$. How about $n \in \mathbb{N}$. For the value of T_l to be equal to exactly n, two things must happen:

   a) $X_n = l$ (the level l must actually be hit at time n), and
   b) $X_{n-1} \neq l$, $X_{n-2} \neq l$, ..., $X_1 \neq l$, $X_0 \neq l$ (the level l has not been hit before).

   Therefore,

$$G^n(x_0, x_1, \dots, x_n) = \begin{cases} 1, & x_0 \neq l, x_1 \neq l, \dots, x_{n-1} \neq l, x_n = l \\ 0, & \text{otherwise.} \end{cases}$$

The hitting time T_2 of the level $l = 2$ for a particular trajectory of a symmetric simple random walk is depicted below:

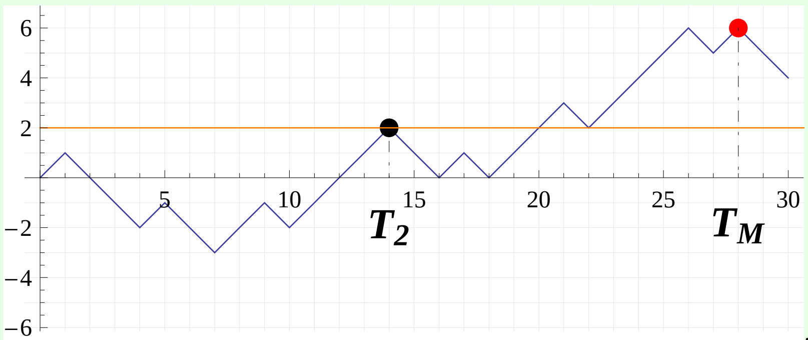

.

3. How about something that is *not* a stopping time? Let $T \in \mathbb{N}$ be an arbitrary time-horizon and let T_M be the last time during $0, \dots, T$ that the random walk visits its maximum during $0, \dots, T$ (see picture above, where $T = 30$). If you bought a stock at time $n = 0$, had to sell it some time before or at T and had the ability to predict the future, this is one of the points you would choose to sell it at. Of course, it is impossible in general to decide whether $T_M = n$, for some $n \in 0, \dots, T-1$ without the knowledge of the values of the random walk after n. More precisely, let us sketch the proof of the fact that T_M is not a stopping time. Suppose, to the contrary, that it is, and let G^n be the family of decision functions. Consider the following two trajectories: $(0, 1, 2, 3, \dots, T-1, T)$ and $(0, 1, 2, 3, \dots, T-1, T-2)$. They differ only in the direction of the last step. They also differ in the fact that $T_M = T$ for the first one and $T_M = T-1$ for the second one. On the other hand, by the definition of the decision functions, we have

$$\mathbf{1}_{\{T_M = T-1\}} = G^{T-1}(X_0, \dots, X_{T-1}).$$

The right-hand side is equal for both trajectories, while the left-hand side equals to 0 for the first one and 1 for the second one. A contradiction.

5.2 Wald's identity II and Gambler's Ruin

Having defined the notion of a stopping time, let us try to compute something about it. The random variables $\{\xi_n\}_{n \in \mathbb{N}}$ in the statement of the theorem below are only assumed to be independent of each other and identically distributed. To make things simpler, you can think of $\{\xi_n\}_{n \in \mathbb{N}}$ as increments of a simple random walk. Before we state the main result, here is an extremely useful identity:

Proposition 5.2.1 (Tail formula for the expectation)**.** *Let N be an $\mathbb{N}_0$-valued random variable. Then*

$$\mathbb{E}[N] = \sum_{k=1}^{\infty} \mathbb{P}[N \geq k].$$

Proof. Clearly, $\mathbb{P}[N \geq k] = \sum_{j \geq k} \mathbb{P}[N = j]$, so (note what happens to the indices when we switch the sums)

$$\sum_{k=1}^{\infty} \mathbb{P}[N \geq k] = \sum_{k=1}^{\infty} \sum_{j \geq k} \mathbb{P}[N = j] = \sum_{j=1}^{\infty} \sum_{k=1}^{j} \mathbb{P}[n = j] = \sum_{j=1}^{\infty} j\mathbb{P}[N = j] = \mathbb{E}[N].$$

□

Theorem 5.2.2 (Wald's Identity II)**.** *Let $\{\xi_n\}_{n\in\mathbb{N}}$ be a sequence of independent, identically distributed random variables with $\mathbb{E}[|\xi_1|] < \infty$. Set*

$$X_n = \sum_{k=1}^{n} \xi_k,\ n \in \mathbb{N}_0.$$

If T is an $\{X_n\}_{n\in\mathbb{N}_0}$-stopping time such that $\mathbb{E}[T] < \infty$, then

$$\mathbb{E}[X_T] = \mathbb{E}[\xi_1]\mathbb{E}[T].$$

Proof. Here is another way of writing the sum $\sum_{k=1}^{T} \xi_k$:

$$\sum_{k=1}^{T} \xi_k = \sum_{k=1}^{\infty} \xi_k \mathbf{1}_{\{k \leq T\}}.$$

The idea behind it is simple: add all the values of ξ_k for $k \leq T$ and keep adding zeros (since $\xi_k \mathbf{1}_{\{k\leq T\}} = 0$ for $k > T$) after that. Taking expectation of both sides and switching $\mathbb{E}$ and $\sum$ (this can be justified, but the argument is technical and we omit it here) yields:

$$\mathbb{E}[\sum_{k=1}^{T} \xi_k] = \sum_{k=1}^{\infty} \mathbb{E}[\mathbf{1}_{\{k\leq T\}} \xi_k]. \tag{5.2.1}$$

Let us examine the term $\mathbb{E}[\xi_k \mathbf{1}_{\{k\leq T\}}]$ in some detail. We first note that

$$\mathbf{1}_{\{k\leq T\}} = 1 - \mathbf{1}_{\{k>T\}} = 1 - \mathbf{1}_{\{k-1\geq T\}} = 1 - \sum_{j=0}^{k-1} \mathbf{1}_{\{T=j\}}.$$

Therefore,

$$\mathbb{E}[\xi_k \mathbf{1}_{\{k \le T\}}] = \mathbb{E}[\xi_k] - \sum_{j=0}^{k-1} \mathbb{E}[\xi_k \mathbf{1}_{\{T=j\}}].$$

By the assumption that T is a stopping time, the indicator $\mathbf{1}_{\{T=j\}}$ can be represented as $\mathbf{1}_{\{T=j\}} = G^j(X_0, \dots, X_j)$, and, because each X_i is just a sum of the increments, we can actually write $\mathbf{1}_{\{T=j\}}$ as a function of $\xi_1, \dots, \xi_j$ only - say $\mathbf{1}_{\{T=j\}} = H^j(\xi_1, \dots, \xi_j)$. By the independence of $(\xi_1, \dots, \xi_j)$ from ξ_k (because $j < k$) we have

$$\begin{aligned}\mathbb{E}[\xi_k \mathbf{1}_{\{T=j\}}] &= \mathbb{E}[\xi_k H^j(\xi_1, \dots, \xi_j)] = \mathbb{E}[\xi_k]\mathbb{E}[H^j(\xi_1, \dots, \xi_j)] \\ &= \mathbb{E}[\xi_k]\mathbb{E}[\mathbf{1}_{\{T=j\}}] = \mathbb{E}[\xi_k]\mathbb{P}[T = j].\end{aligned}$$

Therefore,

$$\begin{aligned}\mathbb{E}[\xi_k \mathbf{1}_{\{k \le T\}}] &= \mathbb{E}[\xi_k] - \sum_{j=0}^{k-1} \mathbb{E}[\xi_k]\mathbb{P}[T = j] = \mathbb{E}[\xi_k]\mathbb{P}[T \ge k] \\ &= \mathbb{E}[\xi_1]\mathbb{P}[T \ge k],\end{aligned}$$

where the last equality follows from the fact that all ξ_k have the same distribution.

Going back to (5.2.1), we get

$$\begin{aligned}\mathbb{E}[X_T] &= \mathbb{E}[\sum_{k=1}^{T} \xi_k] = \sum_{k=1}^{\infty} \mathbb{E}[\xi_1]\mathbb{P}[T \ge k] \\ &= \mathbb{E}[\xi_1] \sum_{k=1}^{\infty} \mathbb{P}[T \ge k] = \mathbb{E}[\xi_1]\mathbb{E}[T],\end{aligned}$$

where we use Proposition 5.2.1 for the last equality. □

Example 5.2.3 (Gambler's ruin problem)**.** . A gambler starts with $\$x$ dollars and repeatedly plays a game in which he wins a dollar with probability $\frac{1}{2}$ and loses a dollar with probability $\frac{1}{2}$. He decides to stop when one of the following two things happens:

1. he goes bankrupt, i.e., his wealth hits 0, or
2. he makes enough money, i.e., his wealth reaches some level $a > x$.

The classical "Gambler's ruin" (dating at least to 1600s) problem asks the following question: what is the probability that the gambler will make a dollars before he goes bankrupt?

Gambler's wealth $\{W_n\}_{n \in \mathbb{N}}$ is modeled by a simple random walk starting from x, whose increments $\xi_k = W_k - W_{k-1}$ are coin-tosses.

Then $W_n = x + X_n$, where $X_n = \sum_{k=0}^{n} \xi_k$, $n \in \mathbb{N}_0$. Let T be the time the gambler stops. We can represent T in two different (but equivalent) ways. On the one hand, we can think of T as the smaller of the two hitting times T_{-x} and T_{a-x} of the levels $-x$ and $a-x$ for the random walk $\{X_n\}_{n\in\mathbb{N}_0}$ (remember that $W_n = x + X_n$, so these two correspond to the hitting times for the process $\{W_n\}_{n\in\mathbb{N}_0}$ of the levels 0 and a). On the other hand, we can think of T as the first hitting time of the two-element *set* $\{-x, a-x\}$ for the process $\{X_n\}_{n\in\mathbb{N}_0}$. In either case, it is quite clear that T is a stopping time (can you write down the decision functions?). When we talked about the maximum of the simple symmetric random walk, we proved that it hits any value if given enough time. Therefore, the probability that that the gambler's wealth will remain strictly between 0 and a forever is zero and so, $\mathbb{P}[T < \infty] = 1$.

What can we say about the random variable X_T - the gambler's wealth (minus x) at the *random* time T? Clearly, it is either equal to $-x$ or to $a-x$, and the probabilities p_0 and p_a with which it takes these values are exactly what we are after in this problem. We know that, since there are no other values X_T can take, we must have $p_0 + p_a = 1$. Second Wald's identity gives us the second equation for p_0 and p_a:

$$\mathbb{E}[X_T] = \mathbb{E}[\xi_1]\mathbb{E}[T] = 0 \cdot \mathbb{E}[T] = 0,$$

so

$$0 = \mathbb{E}[X_T] = p_0(-x) + p_a(a-x).$$

These two linear equations with two unknowns yield

$$p_0 = \frac{a-x}{a},\ p_a = \frac{x}{a}.$$

It is remarkable that the two probabilities are proportional to the amounts of money the gambler needs to make (lose) in the two outcomes. The situation is different when $p \neq \frac{1}{2}$.

5.3 The distribution of the first hitting time of a simple symmetric random walk

Let $\{X_n\}_{n\in\mathbb{N}_0}$ be a simple random walk, with the probability p of stepping up. Let $T_1 = \min\{n \in \mathbb{N}_0 : X_n = 1\}$ be the first hitting time of level $l = 1$, and let $\{p_n\}_{n\in\mathbb{N}_0}$ be its pmf, i.e., $p_n = \mathbb{P}[T_1 = n]$, $n \in \mathbb{N}_0$. The goal of this section is to use the powerful generating-function methods to find $\{p_n\}_{n\in\mathbb{N}_0}$.

A recursive formula. We start with a simple observation that you cannot get from 0 to 1 for the first time in an even number of steps. Therefore, $p_{2n} = 0$, $n \in \mathbb{N}_0$. Also, $p_1 = p$ - you simply have to go up on the first step.

What about $n > 1$? In order to go from 0 to 1 in $n > 1$ steps (and not before!) the first step needs to be *down* and then you need to climb up from -1 to 1 in $n-1$ steps. Climbing from -1 to 1 is exactly the same as climbing from -1 to 0 and then climbing from 0 to 1. If it took j steps to go from -1 to 0 it will have to take $n-1-j$ steps to go from 1 to 2, where j can be anything from 1 to $n-2$, in order to finish the job in exactly $n-1$ steps. So, using formulas, we have

$$\begin{aligned}\mathbb{P}[T_1 = n] = q \sum_{j=1}^{n-2} \mathbb{P}\Big[&\text{“exactly } j \text{ steps to first hit 0 from } -1\text{”} \\ &\textit{and} \text{ “exactly } n-1-j \text{ steps to first hit 1 from 0”}\Big].\end{aligned} \tag{5.3.1}$$

But there is nothing special about 0 as a starting point. Taking j steps from -1 to 0 is exactly the same as taking j steps from 0 to 1, so

$$\mathbb{P}[\text{“exactly } j \text{ steps to first hit 0 from } -1\text{”}] = \mathbb{P}[T_1 = j] = p_j.$$

By the same token,

$$\begin{aligned}\mathbb{P}[\text{“exactly n-1-j steps to first hit 1 from 0”}] &= \mathbb{P}[T_1 = n-1-j] \\ &= p_{n-1-j}.\end{aligned}$$

Finally, I claim that the two events are independent of each other. Indeed, once we have reached 0, the future increments of the random walk behave exactly the same as the increments of a fresh random walk starting from zero - they are independent of everything that happened in the past. Equivalently, a knowledge of everything that happened until the moment the random walk hit 0 for the first time does not change our perception (and estimation) of what is going to happen later (in this case the likelihood of hitting 1 in exactly $n-1-j$ steps). This property is called the *regeneration property* or the *strong Lévy property* of random walks. More precisely (but still not entirely precise), we can make the following claim:

> Let $\{X_n\}_{n\in\mathbb{N}_0}$ be a simple random walk and let T be any $\mathbb{N}_0$-valued stopping time. Define the process $\{Y_n\}_{n\in\mathbb{N}_0}$ by $Y_n = X_{T+n} - X_T$. Then $\{Y_n\}_{n\in\mathbb{N}_0}$ is also a simple random walk, and it is independent of X up to T.

In order to check your understanding, try to convince yourself that the requirement that T be a stopping time is necessary - find an example of a random time T which is not a stopping time where the statement above fails.

We can go back to the distribution of the hitting time T_1, and use our newly-found independence together with (5.3.1) to obtain the following recursion

$$p_n = q \sum_{j=1}^{n-2} p_j p_{n-j-1}, n > 1, \qquad p_0 = 0,\ p_1 = p. \tag{5.3.2}$$

Generating-function approach This is where generating functions step in. We will use (5.3.2) to derive an equation for the generating function $P(s) = \sum_{k=0}^{\infty} p_k s^k$. The sum on the right-hand side of (5.3.2) looks a little bit like a convolution, so let us compare it to the following expansion of the square $P(s)^2$:

$$P(s)^2 = \sum_{k=0}^{\infty} (\sum_{i=0}^{k} p_i p_{k-i}) s^k.$$

The inner sum $\sum_{i=0}^{k} p_i p_{k-i}$ needs to be split into several parts to get an expression which matches (5.3.2):

$$\begin{aligned}\sum_{i=0}^{k} p_i p_{k-i} &= \underbrace{p_0 p_k}_{0} + \sum_{i=1}^{k-1} p_i p_{k-i} + \underbrace{p_k p_0}_{0} = \sum_{i=1}^{(k+1)-2} p_i p_{(k+1)-i-1} \\ &= q^{-1} p_{k+1}, \text{ for } k \geq 2.\end{aligned}$$

Therefore, since the coefficients of $P(s)^2$ start at s^2, we have

$$qsP(s)^2 = qs \sum_{k=2}^{\infty} q^{-1} p_{k+1} s^k = \sum_{k=2}^{\infty} p_{k+1} s^{k+1} = P(s) - ps,$$

which is nothing but a quadratic equation for P.

Remark 5.3.1. Here is another - shorter, but less rigorous - way of deriving the same equation for P. The first hitting time T_1 can be written as follows:

$$T_1 = \begin{cases} 1, & X_1 = 1, \\ 1 + T_1' + T_1'', & X_1 = -1, \end{cases}$$

where T_1' is the time it takes to hit 0 from -1 and T_1'' is the time it takes to hit 1 from 0. The notation T_1' and T_1'' is suggestive of the fact that T_1' and T_1'' have the same distribution as T_1. Moreover, they are independent of each other, as argued above. Using the expression $P_X(s) = \mathbb{E}[s^X]$ and the law of total probability with conditioning on the first step, we get

$$\begin{aligned} P(s) = \mathbb{E}[s^{T_1}] &= \mathbb{E}[s^{T_1} | X_1 = 1] \times \mathbb{P}[X_1 = 1] + \mathbb{E}[s^{T_2} | X_1 = -1] \times \mathbb{P}[X_1 = -1] \\ &= p\mathbb{E}[s^1 | X_1 = 1] + q\mathbb{E}[s^{1+T_1'+T_1''}] = ps + q\mathbb{E}[s^{1+T_1'+T_1''}] \\ &= ps + qs\mathbb{E}[s^{T_1'}] \times \mathbb{E}[S^{T_1''}] = ps + qsP(s)^2. \end{aligned}$$

Now that we have the following equation for P:

$$P(s) = ps + qsP(s)^2, \tag{5.3.3}$$

the first task is to solve it. It is a quadratic equation in $P(s)$, so so it admits two solutions (for each s):

$$P(s) = \frac{1 \pm \sqrt{1-4pqs^2}}{2qs}.$$

One of the two solutions is always greater than 1 in absolute value, so it cannot correspond to a value of a generating function. Therefore,

$$P(s) = \frac{1 - \sqrt{1-4pqs^2}}{2qs}, \text{ for } |s| \le \frac{1}{2\sqrt{pq}}.$$

It remains to extract the information about $\{p_n\}_{n\in\mathbb{N}_0}$ from P. We will not derive expressions for all p_n (but see the last problem in the Problems section), because we do not need to. We can get very useful information from P itself.

Do we actually hit 1 sooner or later? What happens if we try to evaluate $P(1)$? We should get 1, right? In fact, what we get is the following:

$$P(1) = \frac{1-\sqrt{1-4pq}}{2q} = \frac{1-|p-q|}{2q} = \begin{cases} 1, & p \ge \frac{1}{2} \\ \frac{p}{q}, & p < \frac{1}{2} \end{cases}$$

Clearly, $P(1) < 1$ when $p < q$. The explanation is simple - the random walk may fail to hit the level 1 *at all*, if $p < q$. In that case $P(1) = \sum_{k=0}^{\infty} p_k = \mathbb{P}[T_1 < \infty] < 1$, or, equivalently, $\mathbb{P}[T_1 = +\infty] > 0$. It is remarkable that if $p = \frac{1}{2}$, the random walk *will* always hit 1 sooner or later, but this does not need to happen if $p < \frac{1}{2}$. What we have here is an example of a phenomenon known as *criticality*: many physical systems exhibit qualitatively different behavior depending on whether the value of certain parameter p lies above or below certain *critical value* $p = p_c$.

Expected time until we hit 1? Another question that generating functions can help answer is the following one: how long, on average, do we need to wait before 1 is hit? When $p < \frac{1}{2}$, $\mathbb{P}[T_1 = +\infty] > 0$, so we can immediately conclude that $\mathbb{E}[T_1] = +\infty$, by definition. The case $p \ge \frac{1}{2}$ is more interesting. Following the recipe from the lecture on generating functions, we compute the derivative of $P(s)$ and get

$$P'(s) = \frac{2p}{\sqrt{1-4pqs^2}} - \frac{1-\sqrt{1-4pqs^2}}{2qs^2}.$$

- When $p = \frac{1}{2}$, we get

$$\lim_{s \nearrow 1} P'(s) = \lim_{s \nearrow 1} \Big(\frac{1}{\sqrt{1-s^2}} - \frac{1-\sqrt{1-s^2}}{s^2} \Big) = +\infty,$$

  and conclude that $\mathbb{E}[T_1] = +\infty$.

- For $p > \frac{1}{2}$, the situation is less severe:

$$\lim_{s \nearrow 1} P'(s) = \frac{1}{p-q}.$$

We can summarize the situation in the following table

	$\mathbb{P}[T_1 < \infty]$	$\mathbb{E}[T_1]$
$p < \frac{1}{2}$	$\frac{p}{q}$	$+\infty$
$p = \frac{1}{2}$	1	$+\infty$
$p > \frac{1}{2}$	1	$\frac{1}{p-q}$

5.4 Problems

Problem 5.4.1. Either one of the following 4 random times is not a stopping time for a simple random walk $\{X_n\}_{n \in \mathbb{N}_0}$, or they all are. Choose the one which is not in the first case, or choose (e) if you think they all are.

(a) the first hitting time of the level 4,

(b) the first time n such that $X_n - X_{n-1} \neq X_1$,

(c) the first time the walk hits the level 2 or the first time the walk sinks below -5, whatever happens first,

(d) the second time the walk crosses the level 5 or the third time the walk crosses the level -2, whatever happens last,

(e) none of the above.

Problem 5.4.2. Let T_1 be the first hitting time of the level 1, for a simple symmetric random walk $\{X_n\}_{n \in \mathbb{N}_0}$. Then

(a) $\mathbb{P}[T_1 < \infty] = 1$ and $\mathbb{E}[T_1] < \infty$,

(b) $\mathbb{P}[T_1 < \infty] < 1$ and $\mathbb{E}[T_1] < \infty$,

(c) $\mathbb{P}[T_1 < \infty] = 1$ and $\mathbb{E}[T_1] = \infty$,

(d) $\mathbb{P}[T_1 < \infty] < 1$ and $\mathbb{E}[T_1] = \infty$,

(e) none of the above.

Problem 5.4.3. The generating function $P_{T_{-1}}$ of the first hitting time T_{-1} of the level -1 for the simple biased random walk with $p = \mathbb{P}[X_1 = 1]$ satisfies

(a) $P_{T_{-1}}(s) = ps + qsP_{T_{-1}}(s)^2$,

(b) $P_{T_{-1}}(s) = qs + psP_{T_{-1}}(s)^2$,

(c) $P_{T_{-1}}(s) = ps - qsP_{T_{-1}}(s)^2$,

(d) $P_{T_{-1}}(s) = ps + qsP_{T_{-1}}(s)^{-2}$,

(e) none of the above.

Problem 5.4.4. Let $\{X_n\}_{n\in\mathbb{N}_0}$ be a simple biased random walk with $p = \mathbb{P}[X_1 = 1]$, and Let T_l denote the first hitting time of the level l, and P_{T_l} denote its generating function. Then

(a) $P_{T_2}(s) = ps + qsP_{T_2}(s)^4$,

(b) $P_{T_2}(s) = ps^2 + qsP_{T_2}(s)$,

(c) $P_{T_2}(s) = P_{T_1}(s)$,

(d) $P_{T_2}(s) = P_{T_1}(s)^2$,

(e) none of the above.

Problem 5.4.5. If P_{T_a} denotes the generating function of the first hitting time T_a of the level a for the simple symmetric random walk, then

(a) $P_{T_2}(s) = \frac{1}{2}sP_{T_3}(s) + \frac{1}{2}sP_{T_1}(s)$,

(b) $P_{T_2}(s) = \frac{1}{2}P_{T_3}(s) + \frac{1}{2}P_{T_1}(s)$,

(c) $P_{T_2}(s) = sP_{T_3}(s)P_{T_1}(s)$,

(d) $P_{T_2}(s) = \frac{1}{3}(s + P_{T_3}(s) + \frac{1}{2}P_{T_1}(s))$,

(e) none of the above.

Problem 5.4.6. (*) The purpose of this problem is to derive explicit expressions for the probabilities $p_n = \mathbb{P}[T_1 = n]$, where T_1 is the first hitting time of the level 1 for a simple biased random walk $\{X_n\}_{n \in \mathbb{N}_0}$ with $p = \mathbb{P}[X_1 = 1]$. Our starting point is the expression

$$P(s) = \tfrac{1}{2qs}\left(1 - \sqrt{1 - 4pqs^2}\right)$$

for the generating function $P(s) = P_{T_1}(s)$ derived in the notes.

1. Consider the function $f(x) = \sqrt{1-x}$. Find an expression for the its n-th derivative $f^{(n)}(0)$ at $x = 0$. **Note:** The following notation $n!! = n \times (n-2) \times (n-4) \times \dots$, for the product of all positive integers up to n of the same parity as n, is going to be useful.
2. Use the above to expand $P(s)$ in a power-series expansion around $s = 0$ and write down the expressions for the coefficients p_n, $n \in \mathbb{N}_0$.