6.262: Discrete Stochastic Processes 3/14/11

Lecture 12 : Renewal rewards, stopping trials, and Wald's equality

Outline:

- Review strong law for renewals
- Review of residual life
- Time-averages for renewal rewards
- Stopping trials for stochastic processes
- Wald's equality
- Stop when you're ahead

Theorem: If $\{Z_n;\, n \geq 1\}$ **converges to** α **WP1, (i.e.,** $\Pr\{\omega : \lim_n (Z_n(\omega) - \alpha) = 0\} = 1$**), and** $f(x)$ **is continuous at** α**. Then** $\Pr\{\omega : \lim_n f(Z_n(\omega)) = \alpha\} = 1$**.**

For a renewal process with inter-renewals X_i, $0 < \overline{X} < \infty$, $\Pr\Big\{\omega : \lim_n (\frac{1}{n} S_n(\omega) - \overline{X}) = 0\Big\} = 1$**)**

$$\Pr\left\{\omega : \lim_{n\to\infty} \frac{n}{S_n(\omega)} = \frac{1}{\overline{X}}\right\} = 1.$$

For renewal processes, n/S_n **and** $N(t)/t$ **are related by**

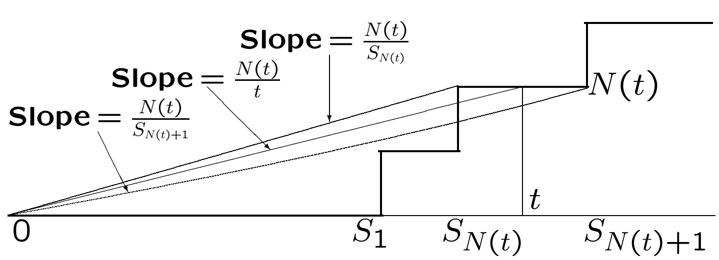

The strong law for renewal processes follows from this relation between n/S_n and $N(t)/t$.

Theorem: For a renewal process with $\overline{X} < \infty$,

$$\Pr\left\{\omega : \lim_{t\to\infty} N(t,\omega)/t = 1/\overline{X}\right\} = 1.$$

This says that the rate of renewals over the infinite time horizon (i.e., $\lim_t N(t)/t$) is $1/\overline{X}$ WP1.

This also implies the weak law for renewals,

$$\lim_{t\to\infty} \Pr\left\{\left|\frac{N(t)}{t} - \frac{1}{\overline{X}}\right| > \epsilon\right\} = 0 \qquad \text{for all } \epsilon > 0$$

Review of residual life

Def: The residual life $Y(t)$ of a renewal process at time t is the remaining time until the next renewal, i.e., $Y(t) = S_{N(t)+1} - t$.

Residual life is a random process; for each sample point ω, $Y(t,\omega)$ is a sample function.

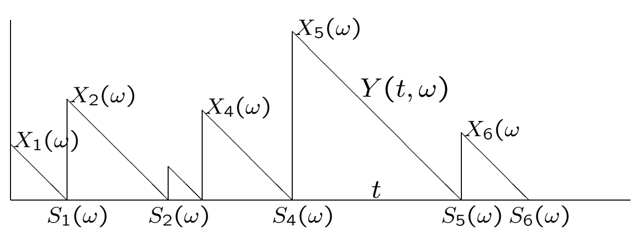

$$\sum_{n=1}^{N(t,\omega)} \frac{X_i^2(\omega)}{2t} \le \frac{1}{t}\int_0^t Y(t,\omega)\,dt \le \sum_{n=1}^{N(t,\omega)+1} \frac{X_i^2(\omega)}{2t}$$

$$\sum_{i=1}^{N(t,\omega)} \frac{X_i^2(\omega)}{2t} \;\le\; \frac{1}{t}\int_0^t Y(t,\omega)\,dt \;\le\; \sum_{i=1}^{N(t,\omega)+1} \frac{X_i^2(\omega)}{2t}$$

Going to the limit $t \to \infty$

$$\begin{aligned} \lim_{t\to\infty} \frac{1}{t}\int_0^t Y(t,\omega)\,dt &= \lim_{t\to\infty} \sum_{n=1}^{N(t,\omega)} \frac{X_i^2(\omega)}{2N(t,\omega)} \frac{N(t,\omega)}{t} \\ &= \frac{\mathsf{E}\left[X^2\right]}{2\mathsf{E}\left[X\right]} \end{aligned}$$

This is infinite if $\mathsf{E}\left[X^2\right] = \infty$. **Think of example where** $\mathsf{p}_X(\epsilon) = 1-\epsilon,\ \mathsf{p}_X(1/\epsilon) = \epsilon$.

Similar examples: Age $Z(t) = t - S_{N(t)}$ **and duration,** $\widetilde{X}(t) = S_{N(t)+1} - S_{N(t)}$.

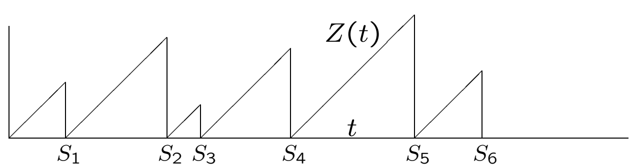

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t \widetilde{X}(\tau)\,d\tau = \frac{\mathsf{E}\left[X^2\right]}{2\mathsf{E}\left[X\right]} \qquad \textbf{WP1}.$$

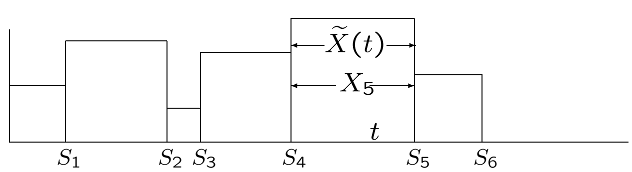

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t \widetilde{X}(\tau)\,d\tau = \frac{\mathsf{E}\left[X^2\right]}{\mathsf{E}\left[X\right]} \qquad \textbf{WP1}.$$

Time-averages for renewal rewards

Residual life, age, and duration are examples of assigning rewards to renewal processes.

The reward $R(t)$ at any time t is restricted to be a function of the inter-renewal period containing t.

In simplest form, $R(t)$ is restricted to be a function $\mathcal{R}(Z(t), \widetilde{X}(t))$.

The time-average for a sample path of $R(t)$ is found by analogy to residual life. Start with the nth inter-renewal interval.

$$R_n(\omega) = \int_{S_{n-1}(\omega)}^{S_n(\omega)} R(t,\omega)\,dt$$

Interval 1 goes from 0 to S_1, with $Z(t) = t$. For interval n, $Z(t) = t - S_{n-1}$, i.e., $S_{N(t)} = S_{n-1}$.

$$\begin{aligned} R_n &= \int_{S_{n-1}}^{S_n} R(t)\,dt \\ &= \int_{S_{n-1}}^{S_n} \mathcal{R}(Z(t), \widetilde{X}(t))\,dt \\ &= \int_{S_{n-1}}^{S_n} \mathcal{R}(t - S_{n-1}, X_n)\,dt \\ &= \int_0^{X_n} \mathcal{R}(z, X_n)\,dz \end{aligned}$$

This is a function only of the rv X_n. Thus

$$\mathsf{E}\left[\mathsf{R}_n\right] = \int_{x=0}^{\infty} \int_{z=0}^{x} \mathcal{R}(z,x)\,dz\,d\mathsf{F}_X(x).$$

Assuming that this expectation exists,

$$\lim_{t\to\infty} \frac{1}{t}\int_{\tau=0}^{t} R(\tau)\,d\tau = \frac{\mathsf{E}\left[\mathsf{R}_n\right]}{\overline{X}} \qquad \textbf{WP1}$$

Example: Suppose we want to find the kth moment of the age.

Then $\mathcal{R}(Z(t), \widetilde{X}(t)) = Z^k(t)$. Thus

$$\begin{aligned}\mathsf{E}[R_n] &= \int_{x=0}^{\infty} \int_{z=0}^{x} z^k \, dz \, d\mathsf{F}_X(x) \\ &= \int_0^{\infty} \frac{x^{k+1}}{k+1} d\mathsf{F}_X(x) = \frac{1}{k} \mathsf{E}\left[X^{k+1}\right]\end{aligned}$$

$$\lim_{t\to\infty} \frac{1}{t} \int_{\tau=0}^{t} R(\tau)\, d\tau = \frac{\mathsf{E}\left[X^{k+1}\right]}{(k+1)\overline{X}} \qquad \textbf{WP1}$$

Stopping trials for stochastic processes

It is often important to analyze the initial segment of a stochastic process, but rather than investigating the interval $(0, t]$ for a fixed t, we want to investigate $(0, t]$ where t is selected by the sample path up until t.

It is somewhat tricky to formalize this, since t becomes a rv which is a function of $\{X(t); \tau \le t\}$. This approach seems circular, so we have to be careful.

We consider only discrete-time processes $\{X_i; i \ge 1\}$.

Let J be a positive integer rv that describes when a sequence $X_1, X_2, \dots$, is to be stopped.

At trial 1, $X_1(\omega)$ is observed and a decision is made, based on $X_1(\omega)$, whether or not to stop. If we stop, $J(\omega) = 1$

At trial 2 (if $J(\omega) \neq 1$), $X_2(\omega)$ is observed and a decision is made, based on $X_1(\omega)$, $X_2(\omega)$, whether or not to stop. If we stop, $J(\omega) = 2$.

At trial 3 (if $J(\omega) \neq 1, 2$), $X_3(\omega)$ is observed and a decision is made, based on $X_1(\omega), X_2(\omega), X_3(\omega)$, whether or not to stop. If we stop, $J(\omega) = 3$, etc.

At each trial n (if stopping has not yet occurred), X_n is observed and a decision (based on $X_1 \dots, X_n$) is made; if we stop, then $J(\omega) = n$.

Def: A stopping trial (or stopping time) J for $\{X_n; n \geq 1\}$, is a positive integer-valued rv such that for each $n \geq 1$, the indicator rv $\mathbb{I}_{\{J=n\}}$ is a function of $\{X_1, X_2, \dots, X_n\}$.

A possibly defective stopping trial is the same except that J might be defective.

We visualize 'conducting' successive trials $X_1, X_2, \dots$, until some n at which the event $\{J = n\}$ occurs; further trials then cease. It is simpler conceptually to visualize stopping the observation of trials after the stopping trial, but continuing to conduct trials.

Since J is a (possibly defective) rv, the events $\{J = 1\}, \{J = 2\}, \dots$ are disjoint.

Example 1: A gambler goes to a casino and gambles until broke.

Example 2: Flip a coin until 10 successive heads appear.

Example 3: Test an hypothesis with repeated trials until one or the other hypothesis is sufficiently probable a posteriori.

Example 4: Observe successive renewals in a renewal process until $S_n \geq 100$.

Suppose the rv's X_i in a process $\{X_n; n \geq 1$ have a finite number of possible sample values. Then any (possibly defective) stopping trial J can be represented as a rooted tree where the trial at which each sample path stops is represented by a terminal node.

Example: X is binary and stopping occurs when the pattern (1, 0) first occurs.

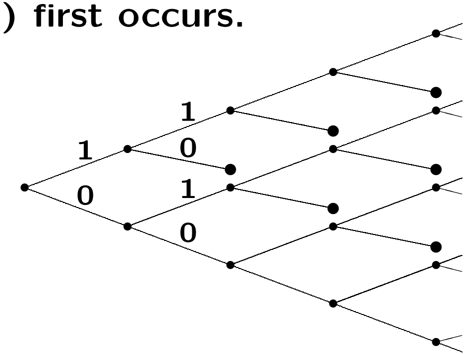

Wald's equality

Theorem (Wald's equality) Let $\{X_n; n \geq 1\}$ be a sequence of IID rv's, each of mean $\overline{X}$. If J is a stopping trial for $\{X_n;\, n \geq 1\}$ and if $\mathsf{E}[J] < \infty$, then the sum $S_J = X_1 + X_2 + \cdots + X_J$ at the stopping trial J satisfies

$$\mathsf{E}[S_J] = \overline{X}\mathsf{E}[J]$$

Prf:

$$S_J = X_1\mathbb{I}_{J\geq 1} + X_2\mathbb{I}_{J\geq 2} + \cdots + X_n\mathbb{I}_{J\geq n} + \cdots$$

$$\mathsf{E}[S_J] = \mathsf{E}\left[\sum_n X_n\mathbb{I}_{J\geq n}\right] = \sum_n \mathsf{E}\left[X_n\mathbb{I}_{J\geq n}\right]$$

The essence of the proof is to show that X_n and $\mathbb{I}_{J\geq n}$ are independent.

To show that X_n and $\mathbb{I}_{J\geq n}$ are independent, note that $\mathbb{I}_{J\geq n} = 1 - \mathbb{I}_{J<n}$. Also $\mathbb{I}_{J<n}$ is a function of $X_1, \ldots, X_{n-1}$. Since the X_i are IID, X_n is independent of $X_1, \ldots, X_{n-1}$, and thus $\mathbb{I}_{J<n}$, and thus of $\mathbb{I}_{J\geq n}$.

This is surprising, since X_n is certainly not independent of $\mathbb{I}_{J=n}$, nor of $\mathbb{I}_{J=n+1}$, etc.

The resolution of this 'paradox' is that, given that $J \geq n$ (i.e., that stopping has not occured before trial n), the trial at which stopping occurs depends on X_n, but whether or not $J \geq n$ occurs depends only on $X_1, \ldots, X_{n-1}$.

Now we can finish the proof.

$$
\begin{aligned}
\mathsf{E}[S_J] &= \sum_n \mathsf{E}\left[X_n \mathbb{I}_{J\ge n}\right] \\
&= \sum_n \mathsf{E}[X_n]\,\mathsf{E}\left[\mathbb{I}_{J\ge n}\right] \\
&= \overline{X}\sum_n \mathsf{E}\left[\mathbb{I}_{J\ge n}\right] \\
&= \overline{X}\sum_n \Pr\{J \ge n\} \;=\; \overline{X}\mathsf{E}[J]
\end{aligned}
$$

In many applications, this gives us one equation in two quantities neither of which is known. Frequently, $\mathsf{E}[S_J]$ is easy to find and this solves for $\mathsf{E}[J]$.

The following example shows, among other things, why $\mathsf{E}[J] < \infty$ is required for Wald's equality.

Stop when you're ahead

Consider tossing a coin with probability of heads equal to p. \$1 is bet on each toss and you win on heads, lose on tails. You stop when your winnings reach \$1.

If $p > 1/2$, your winnings (in the absence of stopping) would grow without bound, passing through 1, so J must be a rv. $S_J = 1$ WP1, so $\mathsf{E}[S_J] = 1$. Thus, Wald says that $\mathsf{E}[J] = 1/\overline{X} = \frac{1}{2p-1}$. Let's verify this in another way.

Note that $J = 1$ with probability p. If $J > 1$, i.e., if $S_1 = -1$, then the only way to reach $S_n = 1$ is to go from $S_1 = -1$ to $S_m = 0$ for some m (requiring $\overline{J}$ steps on average); $\overline{J}$ more steps on average then gets to 1. Thus $\overline{J} = 1 + (1-p)2\overline{J} = \frac{1}{2p-1}$.

Next consider $p < 1/2$. It is still possible to win and stop (for example, $J = 1$ with probability p and $J = 3$ with probability $p^2(1-p)$). It is also possible to head South forever.

Let $\theta = \Pr\{J < \infty\}$. Note that $\Pr\{J = 1\} = p$. Given that $J > 1$, i.e., that $S_1 = -1$, the event $\{J < \infty\}$ requires that $S_m - S_1 = 1$ for some m, and then $S_n - S_m = 1$ for some $n > m$. Each of these are independent events of probability θ, so

$$\theta = p + (1-p)\theta^2$$

There are two solutions, $\theta = p/(1-p)$ and $\theta = 1$, which is impossible. Thus J is defective and Wald's equation is inapplicable.

Finally consider $p = 1/2$. In the limit as p approaches 1/2 from below, $\Pr\{J < \infty\} = 1$. We find other more convincing ways to see this later. However, as p approaches 1/2 from above, we see that $\mathsf{E}[J] = \infty$.

Wald's equality does not hold here, since $\mathsf{E}[J] = \infty$, and in fact does not make sense since $\overline{X} = 0$.

However, you make your $1 with probability 1 in a fair game and can continue to repeat the same feat.

It takes an infinite time, however, and requires access to an infinite capital.

6.262 Discrete Stochastic Processes
Spring 2011