1 Introduction to Renewal Theory

Here, we will present some basic results in renewal theory such as the elementary renewal theorem and the inspection paradox (Section 1), and the renewal reward theorem (Section 2). Our emphasis is on sample-path methods. The reader interested in the renewal reward theorem need not read all of Section 1 beforehand, only Sections 1.1-1.3 suffice.

1.1 Point Processes

Definition 1.1 *A simple point process* $\psi = \{t_n : n \geq 1\}$ *is a sequence of points*

$$0 < t_1 < t_2 < \cdots, \tag{1}$$

with $t_n \longrightarrow \infty$ *as* $n \longrightarrow \infty$. *With* $N(0) \stackrel{\text{def}}{=} 0$ *we let* $N(t)$ *denote the number of points that fall in the interval* $(0, t]$, *and* $\{N(t) : t \geq 0\}$ *is called the counting process for* ψ. *Note that* $N(t) = \max\{n : t_n \leq t\}$. *If the* t_n *are random variables then* ψ *is called a random point process. We sometimes allow a point at the origin and define* $t_0 \stackrel{\text{def}}{=} 0$. $X_n = t_n - t_{n-1},\ n \geq 1$ *is called the* n^{th} *interarrival time.*

We view t as time and view t_n as the n^{th} *arrival* time. The word *simple* refers to the fact that we are not allowing more than one arrival to ocurr at the same time (as is stated precisely in (1)). When the t_n are random variables, ψ is called a random point process.

Note that

$$t_n = X_1 + \cdots + X_n,\ n \geq 1.$$

1.2 Renewal process

A random point process $\psi = \{t_n\}$ for which the interarrival times $\{X_n\}$ form an i.i.d. sequence is called a *renewal process*. t_n is then called the n^{th} *renewal epoch* and $F(x) = P(X \leq x)$ denotes the common interarrival time distribution. The *rate* of the renewal process is defined as $\lambda \stackrel{\text{def}}{=} 1/E(X)$ which is justified by

Theorem 1.1 (Elementary Renewal Theorem (ERT)) *For a renewal process,*

$$\lim_{t\to\infty} \frac{N(t)}{t} = \lambda \text{ w.p.1.}$$

and

$$\lim_{t\to\infty} \frac{E(N(t))}{t} = \lambda.$$

Proof : Observing that $t_{N(t)} \leq t < t_{N(t)+1}$ and that $t_{N(t)} = X_1 + \cdots X_{N(t)}$, yields after division by $N(t)$:

$$\frac{1}{N(t)} \sum_{j=1}^{N(t)} X_j \leq \frac{t}{N(t)} \leq \frac{1}{N(t)} \sum_{j=1}^{N(t)+1} X_j.$$

By the Strong Law of Large Numbers (SLLN), both the left and the right pieces converge to $E(X)$ as $t\longrightarrow\infty$. Since $t/N(t)$ is sandwiched between the two, it also converges to $E(X)$, yielding the first result after taking recipricals.

For the second result, we must show that the collection of rvs $\{N(t)/t : t \geq 1\}$ is uniformly integrable (UI)[1], so as to justify the interchange of limit and expected value,

$$\lim_{t\to\infty} \frac{E(N(t))}{t} = E\Big(\lim_{t\to\infty} \frac{N(t)}{t}\Big).$$

We will show that $P(N(t)/t > x) \leq c/x^2,\ x > 0$ for some $c > 0$ hence proving UI. To this end, choose $a > 0$ such that $0 < F(a) < 1$ (if no such a exists then the renewal process is deterministic and the result is trival). Define new interarrival times via truncation $\hat{X}_n = aI\{X_n > a\}$. Thus $\hat{X}_n = 0$ with probability $F(a)$ and equals a with probability $1 - F(a)$. Letting $\hat{N}(t)$ denote the counting process obtained by using these new interarrival times, it follows that $N(t) \leq \hat{N}(t),\ t \geq 0$. Moreover, arrivals (which now occur in batches) can now only occur at the deterministic lattice of times $\{na : n \geq 0\}$. Letting $p = 1 - F(a)$, and letting K_n denote the number of arrivals that occur at time na, we conclude that $\{K_n\}$ is iid with a geometric distribution with success probability p. Letting $[at]$ denote the smallest integer $\geq at$, we have the inequality

$$N(t) \leq \hat{N}(t) \leq S(t) = \sum_{n=1}^{[at]} K_n,\ t \geq 0.$$

Observing that $E(S(t)) = [at]E(K)$ and $Var(S(t)) = [at]Var(K)$, we obtain $E(S(t)^2) = Var(S(t) + E(S(t))^2 = [at]Var(K) + [at]^2E^2(K) \leq c_1t + c_2t^2$, for constants $c_1 > 0,\ c_2 > 0$. Finally, from Chebychev's inequality, $P(N(t)/t > x) \leq E(N^2(t))/t^2x^2 \leq E(S^2(t))/t^2x^2 \leq c/x^2$ where $c = c_1 + c_2$. ∎

1.3 Forward recurrence time

Since $t_{N(t)} \leq t < t_{N(t)+1}$, we define the *forward recurrence time* as the time until the next point strictly after time t:

$$A(t) \overset{\text{def}}{=} t_{N(t)+1} - t,\ t \geq 0. \tag{2}$$

$A(t)$ is also called the *excess* at time t.

If the arrival times $\{t_n\}$ denote the times at which subways arrive to a platform, then $A(t)$ is the amount of time you must wait for the next subway if you arrive at the platform at time t. If $\{t_n\}$ is a Poisson process at rate λ, then by the memoryless property of the exponential distribution, we know that $A(t) \sim exp(\lambda),\ t \geq 0$. But for a general renewal process, the distribution of $A(t)$ is complicated and depends on the time t.

But by taking the limit as $t\longrightarrow\infty$, we can derive nice formulas for the mean and limiting distribution of $A(t)$:

Proposition 1.1

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t A(s)ds = \frac{E(X^2)}{2E(X)} \text{ w.p.1.}$$

[1] A collection of rvs $\{X_t : t \in T\}$ is said to be uniformly integrable (UI), if $\sup_{t\in T} E(|X_t|I\{|X_t| > x\}) \to 0$, as $x \to \infty$.

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t I\{A(s) > x\}ds = \lambda E(X-x)^+ \text{ w.p.1.}$$

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t P(A(s) > x)ds = \lambda E(X-x)^+,$$

where $a^+ \stackrel{\text{def}}{=} \max\{0, a\}$ *(called the positive part of* a*).*

Proof : The proof is very easy and is based on the same kind of sample-path analysis as used in proving the Elementary Renewal Theorem. For example, for the first statement, we need to compute the area under the graph of $A(s)$, from 0 to t; the sum of the area of triangles with sides X_j. Each triangle has area $X_j^2/2$, and there are $N(t)$ of them by time t; thus,

$$\frac{1}{t}\int_0^t A(s)ds \approx \frac{1}{t}\sum_{j=1}^{N(t)} \frac{X_j^2}{2}.$$

To be precise, we have upper and lower bounds

$$\frac{1}{t}\sum_{j=1}^{N(t)} \frac{X_j^2}{2} \le \frac{1}{t}\int_0^t A(s)ds \le \frac{1}{t}\sum_{j=1}^{N(t)+1} \frac{X_j^2}{2}.$$

Re-writing

$$\frac{1}{t} = \frac{N(t)}{t} \times \frac{1}{N(t)},$$

and using the Elementary Renewal Theorem, and the Strong Law of Large Numbers, then yields the first statement in the Proposition. The second statement is similar with

$$\frac{1}{t}\int_0^t I\{A(s) > x\}ds \approx \frac{1}{t}\sum_{j=1}^{N(t)} (X_j - x)^+;$$

the length of time during the j^{th} interarrival time that $A(s) > x$ is precisely $(X_j - s)^+$. The third statement is obtained by taking the expected value of the second (which can be mathematically justified since the integrand is bounded and non-negative). ∎

1.4 Equilibrium distribution

In terms of the tail of X, $\overline{F}(x) = 1 - F(x) = P(X > x)$, we can compute

$$\lambda E(X-x)^+ = \lambda \int_x^\infty \overline{F}(x)dx, \tag{3}$$

by integrating (over y) the tail $P((X-x)^+ > y) = P(X > x + y)$ to compute the expected value.

As $x \ge 0$ varies, this defines the tail of a distribution, the cdf of which we denote by F_e and call it the *equilibrium distribution* of F:

$$F_e(x) \stackrel{\text{def}}{=} \lambda \int_0^x \overline{F}(x)dx. \tag{4}$$

In the subway example, this then gives us the probability distribution for your waiting time. (The mean of F_e yields average waiting time which is the first part of Proposition 1.1.)

In general for $X \sim F$, we let X_e denote a r.v. with distribution F_e: $P(X_e \le x) = F_e(x)$.

Note that X_e is always a continuous random variable because the density function

$$f_e(x) \stackrel{\text{def}}{=} \frac{d}{dx}F_e(x) = \lambda\overline{F}(x),$$

always exists (even if X is not continuous).

For example, if $P(X = c) = 1$, that is, X is a constant c, then F_e is the uniform distribution on $(0, c)$ as derived as follows: $\lambda = 1/c$, $\overline{F}(x) = 1,\ x \in (0, c)$ and so

$$f_e(x) = \lambda\overline{F}(x) = \frac{1}{c},\ x \in (0, c),$$

the uniform density on $(0, c)$.

As another nice example of forward recurrence time, consider lightbulbs with i.i.d. lifetimes $\{X_n\}$, and a single lamp that is always turned on operating as follows: At time 0 the initial bulb is placed in the lamp with lifetime X_1, and when the bulb burns out at time $t_1 = X_1$, it is instantaneously replaced by the second bulb with lifetime X_2 which burns out at time $t_2 = X_1 + X_2$ and so on. In general, the n^{th} bulb is placed in the lamp at time $t_{n-1} = X_1 + \cdots X_{n-1}$ and burns out at time $t_{n-1} + X_n,\ n \ge 1$.

$A(t)$ represents the remaining lifetime of the bulb that is in the lamp at time t. If you arrive at some time randomly way out in the future, then the distribution of the remaining lifetime of the bulb you find burning is the equilibrium distribution F_e.

1.5 Backwards recurrence time

Since $t_{N(t)} \le t < t_{N(t)+1}$, we define the *backwards recurrence time* as the time since the last point before or at time t:

$$B(t) \stackrel{\text{def}}{=} t - t_{N(t)},\ t \ge 0,$$

where by convention $B(0) \stackrel{\text{def}}{=} 0$. $B(t)$ is also called the *age* at time t, because in the lightbulb lifetime example of the previous section, it represents the age of the light bulb you find burning at time t, namely, how long the bulb has already been burning. In the subway example, if you arrive at time t to the platform, then $B(t)$ represents how long it has been since the last subway arrived.

Analogous to $A(t)$,

Proposition 1.2

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t B(s)ds = \frac{E(X^2)}{2E(X)} \text{ w.p.1.}$$

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t I\{B(s) > x\}ds = \lambda E(X - x)^+ \text{ w.p.1.}$$

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t P(B(s) > x)ds = \lambda E(X - x)^+.$$

And so, the stationary distribution of $B(t)$ is the same as that of $A(t)$; $F_e(x)$. This should not be surprising because the graph of $B(t)$ is just the mirror image of that of $A(t)$, and so the following are valid as they were for $A(t)$:

$$\frac{1}{t}\int_0^t B(s)ds \approx \frac{1}{t}\sum_{j=1}^{N(t)} \frac{X_j^2}{2}.$$

$$\frac{1}{t}\int_0^t I\{B(s) > x\}ds \approx \frac{1}{t}\sum_{j=1}^{N(t)} (X_j - x)^+.$$

In general for $X \sim F$, we let X_b denote a r.v. with the stationary distribution for $B(t)$; $P(X_b \le x) = F_e(x)$.

1.6 Spread

Since $t_{N(t)} \le t < t_{N(t)+1}$, we define the *spread* as the length of the interarrival time containing t:

$$S(t) \stackrel{\text{def}}{=} t_{N(t)+1} - t_{N(t)} = B(t) + A(t),\ t \ge 0.$$

If $N(t) = n$ then $S(t) = X_{n+1}$, so that in general

$$S(t) = X_{N(t)+1}.$$

Since $S(t) = B(t) + A(t)$, It follows immediately from the previous two propositions that

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t S(s)ds = \frac{E(X^2)}{E(X)} \text{ w.p.1.}$$

In terms of the lightbulb example, this means that if you randomly (way out in the future) observe the lamp then the lightbulb you find burning has a (total) lifetime with mean $E(X^2)/E(X)$, not $E(X)$! In fact since $E(X^2) \ge (E(X))^2$, we conclude that

$$\frac{E(X^2)}{E(X)} \ge E(X),$$

that is, that the mean is larger than that of a typical lifetime. This is part of a more general result called the *Inspection Paradox* which we will address in the next section. First we consider the stationary distribution of spread:

Proposition 1.3

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t I\{S(s) > x\}ds = \lambda E(XI\{X > x\}) \text{ w.p.1.}$$

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t P(S(s) > x)ds = \lambda E(XI\{X > x\}).$$

Integrating yields

$$\lambda E(XI\{X > x\}) = \lambda x\overline{F}(x) + \overline{F}_e(x).$$

The proof is based on

$$\frac{1}{t}\int_0^t I\{S(s)>x\}ds \approx \frac{1}{t}\sum_{j=1}^{N(t)} X_j I\{X_j>x\}.$$

The length of time during the j^{th} interarrival time that $S(s)>x$ is precisely X_j if $X_j>x$; 0 otherwise.

We define the tail

$$\overline{F}_s(x)=\lambda x\overline{F}(x)+\overline{F}_e(x),$$

and let X_s denote a r.v. with this distribution: $P(X_s\le x)=F_s(x)$. Note that if F has a density f then (via taking derivatives) F_s has density $f_s(x)=\lambda x f(x)$. Unlike X_e however, X_s is not a continuous r.v. in general. For example if $P(X=c)=1$ then $S(t)=c,\ t\ge 0$ and so $P(X_s=c)=1$; $F_s=F$ in this case.

We summarize with:

If F has a density f then F_s has a density given by

$$f_s(x)=\lambda x f(x).$$

Finally observe that in the density case,

$$\begin{aligned} E(X_s) &= \int_0^\infty x f_s(x)dx \\ &= \int_0^\infty \lambda x^2 f(x)dx \\ &= \lambda E(X^2) \\ &= \frac{E(X^2)}{E(X)}, \end{aligned}$$

which agrees with our time average derivation, $\lim_{t\longrightarrow\infty}\frac{1}{t}\int_0^t S(s)ds$, above. In general (density case or not), one can integrate the tail $\overline{F}_s(x)=\lambda x\overline{F}(x)+\overline{F}_e(x)$, and get the same answer.

1.7 Inspection paradox

We already proved that $E(X_s)\ge E(X)$ and we now extend this to

Proposition 1.4 (Inspection Paradox) *For every fixed $t\ge 0$, $S(t)$ is stochastically larger than X, that is,*

$$P(S(t)>x)\ge P(X>x),\ t\ge 0,\ x\ge 0.$$

Moreover X_s is stochastically larger that X:

$$P(X_s>x)\ge P(X>x),\ x\ge 0.$$

Note that in terms of the distributions F and F_s we can rewrite the above as

$$P(S(t)>x)\ge \overline{F}(x),\ t\ge 0,\ x\ge 0,$$

and

$$P(X_s > x) \geq \overline{F}(x), \; x \geq 0.$$

Since expected values can be computed by integrating the tails, we get

$$E(S(t)) = \int_0^\infty P(S(t) > x)dx \geq \int_0^\infty P(X > x)dx = E(X),$$
$$E(X_s) = \int_0^\infty P(X_s > x)dx \geq \int_0^\infty P(X > x)dx = E(X),$$

that is, $E(S(t)) \geq E(X)$ and $E(X_s) \geq E(X)$, the second of which we already proved in the last section.
Proof :[of the Inspection Paradox]

$$\begin{aligned} P(S(t) > x | N(t) = n, t_n = s) &= P(X_{n+1} > x | X_{n+1} > t - s) \\ &= \frac{P(X_{n+1} > x, X_{n+1} > t - s)}{P(X_{n+1} > t - s)} \\ &= \frac{\overline{F}(\max(x, t - s))}{\overline{F}(t - s)}. \end{aligned}$$

We next will show that

$$\frac{\overline{F}(\max(x, t - s))}{\overline{F}(t - s)} \geq \overline{F}(x),$$

which being a statement that is independent of both n and s yields the desired result. To this end, note that if $x > t - s$ then $\max(x, t - s) = x$ in which case

$$\frac{\overline{F}(\max(x, t - s))}{\overline{F}(t - s)} = \frac{\overline{F}(x)}{\overline{F}(t - s)} \geq \overline{F}(x),$$

because the denominator is always less than or equal to 1. If $x \leq t - s$, then $\max(x, t - s) = t - s$ in which case

$$\frac{\overline{F}(\max(x, t - s))}{\overline{F}(t - s)} = \frac{\overline{F}(t - s)}{\overline{F}(t - s)} = 1 \geq \overline{F}(x).$$

This completes the proof of the first statement. For the second statement we take limits of the first:

$$\begin{aligned} P(X_s \geq x) &= \lim_{t \to \infty} \frac{1}{t} \int_0^t P(S(s) > x)ds \\ &\geq \lim_{t \to \infty} \frac{1}{t} \int_0^t P(X > x)ds \\ &\geq P(X > x). \end{aligned}$$

■

1.8 Implications of the inspection paradox

The lightbulb example contains the essentials of the inspection paradox: The bulb you find burning is biased towards having a longer lifetime than usual. This is because longer lifetimes cover more of the time line. Your observation time is more likely to fall inside a large lifetime. In the case when F has a density $f(x)$ so that F_s has density $f_s(x) = \lambda x f(x)$, this *length biasing* effect is clear: the probability of observing a lifetime in progress of length x is proportional to x.

As an extreme case of the inspection paradox, consider lightbulbs that with probability 0.9 are defective, that is, blow out immediately, and with probability 0.1 live for exactly 1 month: $P(X = 0) = 0.9,\ P(X = 1) = 0.1$. If you observe a burning bulb, then it must be a bulb with a length 1 lifetime; you will never observe a defective one. To be precise, $P(X_s = 1) = 1$.

If an inspector's job is to estimate the true mean lifetime of bulbs that have an apriori unknown distribution F, then the WRONG thing to do is to inspect burning lights at random times way out in the future: by doing so the inspector will observe bulbs with lifetime distribution F_s and obtain an answer that is too large. For example, in the $P(X = 0) = 0.9,\ P(X = 1) = 0.1$ case above, the inspector would conclude that all bulbs have a lifetime of 1! The correct way is to take a large sample (of size n) of new bulbs and determine the lifetime of each one and then average them. By the strong law of large numbers this ensures getting the correct answer:

$$\frac{1}{n}\sum_{j=1}^{n} X_j \approx E(X).$$

It is interesting to see what happens in the case when $X \sim exp(\lambda)$, for then, by the memoryless property, X_e and X_b are i.i.d. $\sim exp(\lambda)$. Since $X_s = X_b + X_e$, we conclude that $X_s \sim gamma(2, \lambda)$. In particular, $E(X_s) = 2E(X)$; it has a mean that is twice as large as a typical one!

ln general, X_b and X_e are dependent random variables. For example, if $P(X = c) = 1$, then $X_b + X_e = c$ and so $X_b = c - X_e$. In this case, both X_b and X_e have a $unif(0, c)$ distribution.

In general, it would be nice to derive the joint distribution of backward and forward recurrence time:

$$P(X_b > y, X_e > x) \stackrel{\text{def}}{=} \lim_{t\to\infty} \frac{1}{t}\int_0^t P(B(s) > y, A(s) > x)ds.$$

Proposition 1.5

$$P(X_b > y, X_e > x) = \overline{F}_e(y + x).$$

The above is proved by observing that

$$\frac{1}{t}\int_0^t I\{B(s) > y, A(s) > x\}ds \approx \frac{1}{t}\sum_{j=1}^{N(t)} (X_j - (y + x))^+\}.$$

The length of time during the j^{th} interarrival time that $B(s) > y$ and $A(s) > x$ is precisely $(X_j - (y + x))^+$. To see this, consider the first interarrival time, so that $B(s) = s$ and $A(s) = X_1 - s,\ s \in [0, X_1)$. Then $I\{B(s) > y, A(s) > x\} = I\{s > y, X_1 - s > x\} = I\{y < s < X_1 - x\}$. If $X_1 - (y + x) > 0$, then this indicator is non-zero and $\int_0^{X_1} I\{y < s < X_1 - x\}ds = \int_y^{X_1-x} ds = X_1 - (y + x)$. If $X_1 - (y + x) \le 0$, then the indicator is 0 and $\int_0^{X_1} I\{y < s < X_1 - x\}ds = 0$. Thus we get $(X_1 - (y + x))^+$.

1.9 Examples

1. Suppose Downtown subways arrive to the West 116th street station exactly every 15 minutes. If you arrive at random to the station way out in the future; how long on average must you wait? **SOLUTION:** We want mean forward recurrence time, $E(X_e) = E(X^2)/2E(X) = (15)^2/2(15) = 7.5$. Here X is a typical subway interarrival time; $P(X = 15) = 1$. Thus $E(X^2) = X^2 = 15^2 = 225,\ E(X) = X = 15$. Note further that this is the mean of the equilibrium distribution F_e which is $Unif(0, 15)$ and hence has mean 7.5. $F_e(x) = x/15,\ x \in (0, 15)$.

2. In the above example, what is the probability that you must wait longer than 5 minutes? **SOLUTION:** $X_e \sim Unif(0, 15)$; $F_e(x) = x/15,\ x \in (0, 15)$. Thus $P(X_e > 5) = \overline{F}_e(5) = 10/15 = 2/3$.

3. Computer monitors lifetimes have mean $E(X) = 3$ years and variance $Var(X) = 2$. Everytime the monitor breaks, immediately a new one is installed. You arrive at random out in the future and find a working monitor. What is its expected total lifetime? **SOLUTION:** We want mean spread, $E(X^2)/E(X)$. Noting that $Var(X) = E(X^2) - E^2(X)$, we obtain $E(X^2) = Var(X) + E^2(X) = 2 + 9 = 11$. Thus$E(X^2)/E(X) = 11/3$. Note that this is larger than $E(X) = 3$; the inspection paradox.

4. Trains arrive to a station according to a renewal process with interarrival time distribution $P(X \le x) = F(x) = 1 - e^{-\sqrt{x}},\ x \ge 0$ (called a *Weibull* distribution.) You arrive at random to the station way out in the future. Find the density function of your waiting time. **SOLUTION:** We want $f_e(x) = \lambda \overline{F}(x)$, the density function of the equilibrium distribution. Here, $\overline{F}(x) = e^{-\sqrt{x}}$, and $\lambda = 1/E(X)$, so we need to compute $E(X)$, the mean of the Weibull distribution. Integrating the tail yields

$$\begin{aligned} E(X) &= \int_0^\infty P(X > x)dx \\ &= \int_0^\infty e^{-\sqrt{x}}dx \\ &= \int_0^\infty 2ue^{-u}du \quad \textit{(change of variables } u = \sqrt{x}\textit{)} \\ &= 2. \end{aligned}$$

   Thus $\lambda = 0.5$ and $f_e(x) = 0.5e^{-\sqrt{x}},\ x \ge 0$.

5. Suppose buses heading downtown arrive in front of Columbia University exactly every 15 minutes. Further suppose that independently, Taxis arrive according to a Poisson process at rate 20 per hour. You arrive at random to go downtown. You decide to take either a taxi or a bus, whichever arrives first. How long on average must you wait? **SOLUTION:** The time until the next bus arrives, denote this by X_e, has the $Unif(0, 15)$ distribution, F_e, for the bus interrival times. The time until the next taxi arrives, denote this by T_e, has an exponential distribution with rate 20 per hour, or 1/3 per minute (T_e has the equilibrium distribution for taxi interrival times, exponential by the memoryless property). Thus your waiting time is given by $W = \min\{X_e, T_e\}$ with X_e and T_e independent. We will compute

$E(W)$ by integrating its tail. Since $W > x$ if and only if both $X_e > x$ and $T_e > x$, and since $P(X_e > 15) = 0$, we conclude that $P(W > x) = 0, x \geq 15$, and

$$P(W > x) = P(X_e > x)P(T_e > x) = e^{-(1/3)x}(1 - x/15),\ x \in [0, 15).$$

Thus

$$\begin{aligned} E(W) &= \int_0^\infty P(W > x)dx \\ &= \int_0^{15} e^{-(1/3)x}(1 - x/15)dx, \end{aligned}$$

computational details left to the reader.

2 The Renewal Reward Theorem

2.1 Main result

Consider a NYC taxi driver who drops off passengers at times $t_n,\ n \geq 1$ forming a renewal process with iid interarrival times $X_n = t_n - t_{n-1},\ n \geq 1$ ($t_0 \stackrel{\text{def}}{=} 0$). Suppose that R_j denotes the cost to the j^{th} passenger for their ride. We view this as a *reward* for the driver. (We are assuming a negligeable amount of time is spent by the driver to find new passengers: as soon as one passenger departs, the next one is found immediately.) Our objective is to compute the long run rate at which the driver earns money from the passengers (amount of money per unit time). Letting

$$R(t) = \sum_{j=1}^{N(t)} R_j$$

denote the total amount collected by time t, where $N(t)$ is the counting process for the renewal process, we wish to compute

$$\lim_{t\to\infty} \frac{R(t)}{t}. \tag{5}$$

We suppose that the pairs of rvs (X_j, R_j) are iid which means that R_j is allowed to depend on the length X_j (the length of the ride) but not on any other lengths (or other R_j). Since we can re-write $R(t)$ as

$$\frac{N(t)}{t} \times \frac{1}{N(t)} \sum_{j=1}^{N(t)} R_j,$$

the Elementary Renewal Theorem (ERT) ($N(t)/t \to \lambda = (E(X))^{-1}$) and the Strong Law of Large Numbers (SLLN)($\frac{1}{n}\sum_{j=1}^{n} R_j \to E(R)$) then give (5) as

$$\lim_{t\to\infty} \frac{R(t)}{t} = \frac{E(R)}{E(X)} \text{ w.p.1.}, \tag{6}$$

where (X, R) denotes a typical "cycle" (X_j, R_j).

In words: the rate at which rewards are earned is equal to the expected reward over a "cycle" divided by an expected "cycle length". In terms of taxi rides this means that the rate at which money is earned is equal to the expected cost per taxi ride divided by the expected length of a taxi ride; an intuitively clear result.

For (6) to hold there is no need for rewards to be collected at the end of a cycle; they could be collected at the beginning or in the middle or continuously throughout, but the total amount collected during cycle length X_j is R_j, and it is earned in the time interval $[t_{j-1}, t_j]$. Moreover, "rewards" need not be non-negative (they could be "costs" incurred as opposed to rewards). A precise statement of a theorem follows:

Theorem 2.1 (Renewal Reward Theorem) *For a positive recurrent renewal process in which a "reward" R_j is earned during cycle length X_j and such that $\{(X_j, R_j) : j \geq 1\}$ is iid with $E|R_j| < \infty$, the long run rate at which rewards are earned is given by*

$$\lim_{t\to\infty} \frac{R(t)}{t} = \frac{E(R)}{E(X)} \text{ w.p.1.,} \tag{7}$$

where (X, R) denotes a typical "cycle" (X_j, R_j). In words: the rate at which rewards are earned is equal to the expected reward over a "cycle" divided by an expected "cycle length".

Moreover, we have an expected value version:

$$\lim_{t\to\infty} \frac{E(R(t))}{t} = \frac{E(R)}{E(X)}. \tag{8}$$

Proof : When rewards are non-negative the proof of (7) is based on a "sandwiching" argument in which the two extreme cases (collect at the end of a cycle vs collect at the beginning of a cycle) serve as lower and upper bound respectively:

$$\frac{1}{t}\sum_{j=1}^{N(t)} R_j \leq \frac{R(t)}{t} \leq \frac{1}{t}\sum_{j=1}^{N(t)+1} R_j.$$

Both these bounds converge to $E(R)/E(X)$ yielding the result. In the case when R_j is not non-negative, one can break R_j into positive and negative parts to complete the proof; $R_j = R_j^+ - R_j^-$ with $R_j^+ = \max\{0, R_j\} \geq 0$ and $R_j^- = -\min\{0, R_j\} \geq 0$. Then $R(t) = R^+(t) - R^-(t)$, where

$$R^+(t) = \sum_{j=1}^{N(t)} R_j^+, \; R^-(t) = \sum_{j=1}^{N(t)} R_j^-.$$

The condition $E|R_j| < \infty$ ensures that both $E(R_j^+) < \infty$ and $E(R_j^-) < \infty$ so that the non-negative proof goes through for each of $R^+(t)$ and $R^-(t)$: $R^+(t)/t \to E(R^+)/E(X)$ and $R^-(t)/t \to E(R^-)/E(X)$. Thus, since $E(R) = E(R^+) - E(R^-)$, the result follows, $R(t)/t \to E(R)/E(X)$.

Proof of (8): Note that

$$|R(t)|/t \le Y(t) = \frac{1}{t} \sum_{j=1}^{N(t)+1} |R_j|.$$

Using (7) on $Y(t)$ (rewards are now the $|R_j|$) yields $Y(t) \to E|R|/E(X) < \infty$. Moreover, $N(t)+1$ is a stopping time with respect to $\{X_j, R_j\}$, so from Wald's equation $E(Y(t)) = E(N(t)/t + 1/t)E(|R|) \to E|R|/E(X)$ (via the expected value version of the elementary renewal theorem); thus $\{Y(t) : t \ge 1\}$ is uniformly integrable (UI)[2] and since $|R(t)|/t \le Y(t)$, $\{R(t)/t : t \ge 1\}$ is UI as well.

■

2.2 Examples

The first three examples below involve derivations that we previously provided using a direct sample-path method but which we now place in the renewal reward framework. As the reader is encouraged to check, the previous direct sample-path method is precisely the method we used to prove the renewal reward theorem. But we also can apply (8) to obtain expected value versions too.

1. *Average forward recurrence time.* Our objective is to compute

$$\lim_{t\to\infty} \frac{1}{t} \int_0^t A(s)ds.$$

Recall that $A(t) \stackrel{\text{def}}{=} t_{N(t)+1} - t,\ t \ge 0$, so that if $t_{j-1} \le t < t_j$ then $N(t) = j-1$ and $A(t) = t_j - t$. In order to place this in the context of renewal reward, we collect rewards continuously at rate $r(t) = A(t)$ at time t. For then

$$R(t) = \int_0^t r(s)ds = \int_0^t A(s)ds,$$

and the total reward over X_j is $R_j = X_j^2/2$ since that indeed is the area under $A(s)$ during the cycle X_j. Formally:

$$R_j = \int_{t_{j-1}}^{t_j} A(s)ds = \int_{t_{j-1}}^{t_j} (t_j - s)ds.$$

Changing variables $u = s - t_{j-1}$ yields the integral as

$$R_j = \int_0^{X_j} (X_j - u)du = \Big[X_j u - u^2/2\Big]_0^{X_j} = X_j^2/2.$$

So our answer from the Renewal Reward Theorem is

$$\frac{E(R)}{E(X)} = \frac{E(X^2)}{2E(X)}.$$

[2] A collection $\{X(t)\}$ of non-negative rvs for which $X(t) \to X$ wp1 with $E(X) < \infty$ is UI if and only if $E(X(t)) \to E(X)$.

It is apparent that for any renewal reward problem, as in the example above, we need only integrate over the first cycle X_1 and compute R_1:

$$R_1 = \int_0^{X_1} r(s)du.$$

This is the easiest cycle to compute over (since the lower limit of integration is 0) thus we use $R = R_1$ and $X = X_1$ as our "typical" cycle. Finally, by applying (8) we also have an expected value version

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t E(A(s))ds = \frac{E(X^2)}{2E(X)}.$$

2. *Average backwards recurrence time.* Here we wish to derive

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t B(s)ds.$$

Defining $r(t) = B(t)$ places us in the framework of the Renewal Reward Theorem, and since $B(s) = s,\ 0 \le s < X_1$

$$R = \int_0^X sds = X^2/2,$$

so that our answer becomes the same as for forward recurrence time,

$$\frac{E(R)}{E(X)} = \frac{E(X^2)}{2E(X)}.$$

By applying (8) we also have an expected value version

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t E(B(s))ds = \frac{E(X^2)}{2E(X)}.$$

3. *Stationary distribution of forward recurrence time.* Here we wish to derive

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t I\{A(s) > x\}ds,\ x \ge 0,$$

which can be viewed, in the lightbulb context, as the stationary remaining lifetime distribution, the probability that the bulb you find burning will live for at least x more time units. For x fixed, defining $r(t) = I\{A(t) > x\}$ places us in the framework of the Renewal Reward Theorem, and

$$R = \int_0^X I\{X - s > x\}ds = (X - x)^+,$$

so that our answer becomes

$$\frac{E(X-x)^+}{E(X)} = \lambda\int_x^\infty P(X > s)ds,$$

yielding $\overline{F}_e(x)$, the tail of the *equilibrium distribution* of $F(x) = P(X \le x)$.

4. *Car replacement problem with "T" policy.* Suppose new cars cost $\$C_1$ and have i.i.d. lifetimes $\{V_j : j \geq 1\}$ with a continuous distribution with cdf $F(x) = P(V \leq x)$ (and tail $\overline{F}(x) = 1 - F(x)$). A car that dies when we own it costs $\$C_2$ to tow away (to the dump), then we buy a new one. Suppose that at time 0 we have a new car and then for fixed number $T > 0$ we decide to use the following "T" policy concerning when to buy a new car from then onwards: If our car is still working after T time units then, we give it away to a friend for free and buy a new one. If however, the car dies before T time units, we must pay the tow charge $\$C_2$ and buy a new one. What is our long run cost when using such a policy?

   Letting the consecutive times at which we buy a new car serve as the beginning of a cycle, we conclude that we have a renewal process with interarrival times $X_j = \min\{V_j, T\}$, and that $R_j = C_1 + C_2 I\{V_j < T\}$ is the cost over the j^{th} cycle. Consequently, from the Renewal Reward Theorem, our rate of cost is $E(R)/E(X)$.

   $E(R)$ is immediately computed as

$$E(R) = C_1 + C_2 P(V < T) = C_1 + C_2 F(T),$$

   where $P(V < T) = P(V \leq T) = F(T)$ because F is assumed a continuous distribution. To compute $E(X)$ we integrate the tail of $X = \min\{V, T\}$: $P(X > x) = P(V > x, T > x) = P(V > x)I\{x < T\}$ because T is a constant. Thus

$$E(X) = \int_0^\infty P(X > x)dx = \int_0^T P(V > x)dx = \int_0^T \overline{F}(x)dx.$$

   Finally

$$\frac{E(R)}{E(X)} = g(T) = \frac{C_1 + C_2 F(T)}{\int_0^T \overline{F}(x)dx}. \tag{9}$$

   Of intrinsic interest is now finding the "optimal" value of T to use, the one that minimizes our cost. Clearly, on the one hand, by choosing a T too large, the car will essentially always break down therbye always costing you the C_2 in addition to the C_1. On the other hand, by choosing a T too small, you will essentially keep giving away good cars and have to buy a new one every T time units; incurring C_1 at a fast rate. Between those two extremes should be a moderate value for T that is best. The general method of determing such a value is to differentiate the above function $g(T)$ with respect to T, set equal to 0 and solve. The solution of course depends upon the specific distribution F in use. Several examples are given as homework exercises. Finally note that $E(X)$ can also be computed by using the density function $f(x)$ of V:

$$E(X) = E(\min\{V, T\}) = \int_0^T x f(x)dx + T\overline{F}(T).$$

5. *Taxi driver revisited.* Suppose for the taxi driver problem we incorporate the fact that the driver must spend time finding new passengers. Let Y_j denote

the amount of time spent finding a j^{th} passenger after the $(j-1)^{th}$ passenger departs. Let L_j denote the length of the j^{th} passengers ride, R_j the cost of this ride. We shall assume that $\{Y_j\}$ are i.i.d. and independent of all else (we could more generally only assume that (L_j, Y_j, R_j) are i.i.d. vectors.) Then cycle lengths are now given by $X_j = L_j + Y_j$ and the long run rate at which the driver earns money is given by

$$\frac{E(R)}{E(L) + E(Y)}.$$