1 Some basic renewal theory: The Renewal Reward Theorem

Here, we will present some basic results in renewal theory such as the elementary renewal theorem, and then the very useful *Renewal Reward Theorem* (RRT). As we shall see, there are many applications of the RRT including understanding the famous *inspection paradox*.

1.1 Renewal process

Recall that a random point process $\psi = \{t_n\}$ for which the (non-negative) interarrival times $X_n = t_n - t_{n-1},\ n \geq 1$, form an i.i.d. sequence is called a *renewal process*. t_n is then called the n^{th} *renewal epoch* and $F(x) = P(X \leq x),\ x \geq 0$, denotes the common interarrival time distribution. $t_n = X_1 + \cdots + X_n$, and $N(t) = \max\{n : t_n \leq t\}$ is the counting process. To avoid trivialities we always assume that $P(X > 0) > 0$, hence ensuring that wp1, $t_n \to \infty$, as $n \to \infty$, and $N(t) \to \infty$ as $t \to \infty$. 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 : (First part only) Since $t_n = X_1 + \cdots + X_n,\ n \geq 1$, and

$$t_{N(t)} \leq t < t_{N(t)+1}, \tag{1}$$

we have $t_{N(t)} = \sum_{j=1}^{N(t)} X_j$, and $t_{N(t)+1} = \sum_{j=1}^{N(t)+1} X_j$, yielding after division of (1) 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 wp1 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 reciprocals.

∎

Remark 1 In the elementary renewal theorem, the case when $\lambda = 0$ (e.g., $E(X) = \infty$) is allowed, in which case the renewal process is said to be *null* recurrent. In the case when $0 < \lambda < \infty$ (e.g., $0 < E(X) < \infty$) the renewal process is said to be *positive* recurrent.

1.2 The Renewal Reward Theorem

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

$$\lim_{t\to\infty} \frac{R(t)}{t} = \textit{ the long-run rate that money is earned.} \tag{2}$$

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 $R(t)/t$ 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 (2) as

$$\lim_{t\to\infty} \frac{R(t)}{t} = \frac{E(R)}{E(X)} \text{ w.p.1.}, \tag{3}$$

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 (3) 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). In this more general case, because t is in the middle of a cycle $t_{N(t)} \leq t < t_{N(t)+1}$, we have
>
> $$R(t) = \sum_{j=1}^{N(t)} R_j + l(t),$$

where $l(t)$ is a partial reward, that is, that part of the reward already cumulated in the current cycle (e.g., during $[t_{N(t)}, t]$), and under conditions in the theorem below, asymptotically is negligible; $l(t)/t \to 0$ as $t \to \infty$. A precise statement of the theorem follows:

Theorem 1.2 (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*

$$\lim_{t\to\infty} \frac{R(t)}{t} = \frac{E(R)}{E(X)} = \lambda E(R) \text{ w.p.1.,} \tag{4}$$

where (X, R) denotes a typical "cycle" (X_j, R_j); $\lambda = \{E(X)\}^{-1}$ is the arrival rate for the renewal process. In words: the rate at which rewards are earned is equal to the expected reward over a "cycle" divided by an expected "cycle length". Also note that since λ is the rate at which renewals occur, and each renewal yields a reward on average in the amount $E(R)$, the formula makes good sense as the rate, from basic principles.

Moreover,

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

Proof : ((4) only) When rewards are non-negative the proof of (4) 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, $R(t)/t \to E(R)/E(X)$.

■

1.3 Examples

It is apparent that for any renewal reward problem, we need only compute the reward over the first cycle length X_1, and get R_1. This is the easiest cycle to compute over, thus we let $R = R_1$ and $X = X_1$ as our "typical" cycle.

1. *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 time $T > 0$ 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

$$\frac{E(R)}{E(X)} = g(T) = \frac{C_1 + C_2 F(T)}{\int_0^T \overline{F}(x)dx}. \qquad (6)$$

   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:

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

since we can break up

$$E(\min\{V, T\}) = E(VI\{V \le T\}) + E(TI\{V > T\}).$$

2. *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} passenger's 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

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

3. *Train dispatching problem:* Suppose that passengers arrive to a train platform according to a renewal process at rate μ. As soon as N passengers arrive, a train departs with all N on board. This process continues over and over. Suppose further that the train company incurs a cost at the constant rate of $\$nc$ per unit time whenever exactly n passengers are waiting, and also incurred a fixed cost of $\$K$ each time a train departs. What is the long-run cost rate for the train company?

   We view the departure of trains as the renewal epochs for our "cycles". Letting s_n denote the passenger arrival times and $\{T_n : n \ge 1\}$ denote the interarrival times of passengers, $T_n = s_n - s_{n-1}$, assumed iid with mean $E(T) = 1/\mu$, our first "cycle length" is given by

$$X = T_1 + \cdots + T_N,$$

   and has mean $E(X) = NE(T) = N/\mu$. The cost over the first cycle is given by

$$R = (0)cT_1 + cT_2 + \ldots + (N-1)cT_N + K,$$

   and hence

$$E(R) = cE(T)(1 + 2 + \cdots + N - 1) = cE(T)N(N-1)/2 + K.$$

   Finally,

$$\frac{E(R)}{E(X)} = g(N) = \frac{cE(T)N(N-1)/2 + K}{NE(T)} = \frac{c(N-1)}{2} + \frac{K\mu}{N}.$$

Now suppose the train company wants to minimize cost by choosing the value of N that does so. Then we need to solve $g'(N) = 0$.

To this end, we get

$$g'(N) = c/2 - \frac{K\mu}{N^2} = 0,$$

or

$$N = \sqrt{\frac{2K\mu}{c}}.$$

N is indeed a minimum (as opposed to a maximum) since the second derivative is always positive: $g''(N) = 2K\mu N^{-3} > 0$. If the solution N is not an integer, then one should choose the 2 closest integers surrounding N, $N_1 < N < N_2$, and check which one yields the lower cost, $g(N_1)$ or $g(N_2)$ and that is the answer. If $g(N_1) = g(N_2)$, then either can be used as the answer. As an example 7.13, from S. Ross, *Introduction to Probability Models*, Academic Press, if we consider $K = 6, \mu = 1,\ c = 2$, then $N = \sqrt{6} \approx 2.45$, but as one can check $g(2) = g(3) = 4$.

1.4 Applications to forward recurrence time (excess), age, and the inspection paradox

Consider a renewal point process $\{t_n : n \geq 1\}$ with iid interarrival times $X_n = t_n - t_{n-1},\ n \geq 1$. 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:

$$A(t) \stackrel{\text{def}}{=} t_{N(t)+1} - t,\ t \geq 0. \tag{7}$$

$A(t)$ is also called the *excess* at time t, or *remaining lifetime*. If $t_{n-1} \leq t < t_n$, then $A(t) = t_n - t \leq X_n$.

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. Similarly, if the X_j denote iid lifetimes of lightbulbs, then $A(t)$ denotes the remaining lifetime of the bulb you find burning 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 $t \longrightarrow \infty$, and using the renewal reward theorem, we can derive a nice formula for *average waiting time*:

Proposition 1.1

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

Proof : Note that we can view the iid X_j as cycle lengths, and $r(t) = A(t)$ as the (continuous) rate at time t at which money is being earned. For then the total reward over the first cycle is

$$R = R_1 = \int_0^{X_1} A(s)ds = \int_0^{X_1} (X_1 - s)ds = \int_0^{X_1} sds = X^2/2;$$

the graph of $A(t)$ over the first cycle is a right triangle with sides X_1. In general $R_j = \int_{t_{j-1}}^{t_j} A(s)ds = X_j^2/2,\ j \geq 1$, and indeed the $\{(X_j, R_j)\}$ are iid vectors as is required for using the renewal reward theorem. Thus

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

This illustrates the point that when using the renewal reward theorem, we can collect "rewards" in any way we desire within each cycle; it is only the total amount R_j over the cycle length X_j that matters in the end. ■

Note how, when the point process is Poisson, $E(X^2) = 2/\lambda^2$ and $E(X) = 1/\lambda$ and hence (as should be): $E(X^2)/2E(X) = 1/\lambda = E(X)$. But in general, they are very different, and in fact if $E(X^2) = \infty$ such as in the case when (for example) $P(X > x) = 1/x^3,\ x \geq 1$, then your average waiting time is infinite!

Age

Similar to remaining lifetime (excess) is *age* (backwards recurrence time):

$$B(t) \stackrel{\text{def}}{=} t - t_{N(t)}\ t \geq 0. \tag{8}$$

$B(t)$ denotes the amount of time since the last subway arrived before time t or in the context of the lightbulbs, how long the bulb found burning at time t has already been burning (its age). If $t_{n-1} \leq t < t_n$, then $B(t) = t - t_{n-1} \leq X_n$.

The graph of $B(t)$ over a cycle X_j is also a right triangle with sides X_j but the mirror image of the one for $A(t)$. Thus its area is still $R_j = X_j^2/2$ and the very same renewal reward argument as for $A(t)$ leads to

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

Total lifetime: The inspection paradox

$S(t) = B(t) + A(t) = t_{N(t)+1} - t_{N(t)}$ denotes the length of the interarrival time covering time t. It is sometimes also called the *spread*. If $t_{j-1} \leq t < t_j$, then $S(t) = X_j$. In the context of light bulbs, it represents the total lifetime of the bulb found burning at time t. Defining the reward rate as $r(t) = S(t)$ we note that over X_j its graph is simply a square with sides X_j and hence area X_j^2:

$R_j = \int_{t_{j-1}}^{t_j} S(s)ds = \int_{t_{j-1}}^{t_j} X_j ds = X_j \int_{t_{j-1}}^{t_j} ds = X_j^2$.
Thus $R_j = X_j^2$ and once again the renewal reward theorem yields

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

Since $\sigma^2 = E|X - E(X)|^2 \geq 0$ and $\sigma^2 = E(X^2) - E^2(X)$, we conclude that $E(X^2) \geq E^2(X)$. Thus

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

This says that the average lifetime of the bulb you find burning (at some time t in the infinite future) is larger than the original mean lifetime $E(X)$! This is what is called the *Inspection Paradox*.

Intuition: Your observation (inspection) time t is more likely to land in a larger interval (interarrival time) than usual because larger intervals take up more of the time line. For example, consider the sequence of interarrival times $\{1, 2, 1, 2, 1, 2, \ldots\}$. You will over time, land in a 2 with probability 2/3 and a 1 with only probability 1/3 because the 2s take up 2/3 of the time line. As a more extreme case, consider lightbulbs with iid lifetimes distributed as $P(X = 0) = 0.99$, $P(X = 1) = 0.01$. This means that 99% of the bulbs are defective and blow out immediately. But you would always observe a 1 in progress, because 0 takes up no time. Hence $\lim_{t\to\infty} \frac{1}{t}\int_0^t S(s)ds = 1$, while $E(X) = (0)(0.99) + (1)(0.01) = 0.01$. The difference is huge.

To use the formula $E(X^2)/E(X)$ on this example: Note that $X^2 = X$ since X is either 0 or 1, hence $E(X^2) = E(X)$ yielding $E(X^2)/E(X) = E(X)/E(X) = 1$.

Even for a Poisson process at rate λ, the inspection paradox yields a large difference: $E(X^2)/E(X) = 2/\lambda = 2E(X)$; we get an interval that is twice as large (on average) than the original $E(X)$. In fact, it can be proved that as $t \to \infty$, $S(t)$ for a Poisson process at rate λ converges in distribution to that of $X_1 + X_2$, the sum of 2 independent exponentials a rate λ; an $Erlang(2, \lambda)$ distribution.

The inspection paradox has an even stronger version which we state (without proof) here:

Proposition 1.2 *For a renewal process: For every fixed $t > 0$,*

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

We say that $S(t)$ is stochastically larger than X, for every $t > 0$. Recalling that we can integrate a tail to obtain an expected value then yields, in particular, that $E(S(t)) \geq E(X), \ t > 0$.

What the above Proposition says is that no matter what time (finite) you choose to inspect, you always land in an interval that is stochastically larger than a typical X.