1 Stopping Times

1.1 Stopping Times: Definition

Given a stochstic process $\mathbf{X} = \{X_n : n \geq 0\}$, we view X_n as representing the state of some system at time n. A *random time* τ is a discrete random variable taking values in the time set $\mathbb{N} = \{0, 1, 2, \ldots\}$. X_τ then denotes the state of the system at the random time τ; if $\tau = n$, then $X_\tau = X_n$. If X_n denotes our total fortune right after the n^{th} gamble, then it would be of interest to consider when (at what time n) to stop gambling. If we think of τ as representing such a time to stop (the first time that your fortune reaches \$1000 for example), it seems reasonable that our decision can only depend on what has happened up to that time; information about the future would not be known and thus could not be used to determine our decision to stop. This is the essence of what is called a *stopping time*, which we shall introduce shortly. First we need to formalize the notion of "what has happened up to that time", and this involves information of the stochastic process in question. The *total information known up to time n* is all the information (events) contained in $\{X_0, \ldots, X_n\}$. The point is that we gain more and more information about the process by observing its values consecutively over time.

Definition 1.1 *Let $\mathbf{X} = \{X_n : n \geq 0\}$ be a stochastic process. A stopping time with respect to $\mathbf{X}$ is a random time such that for each $n \geq 0$, the event $\{\tau = n\}$ is completely determined by (at most) the total information known up to time n, $\{X_0, \ldots, X_n\}$. It it not allowed to additionally require knowing any of the future $\{X_{n+1}, X_{n+2}, \ldots\}$.*

In the context of gambling, a stopping time τ is thus a rule that tells us at what time to stop gambling. Our decision to stop after a given gamble can only depend (at most) on the "information" known at that time (not on future information).

If X_n denotes the price of a stock at time n and τ denotes the time at which we will sell the stock (or buy the stock), then our decision to sell (or buy) the stock at a given time can only depend on the information known at that time (not on future information). The time at which one might exercise an option is yet again another example.

Remark 1.1 *All of this can be defined analogously for a sequence $\{X_1, X_2, \ldots\}$ in which time is strictly positive; $n = 1, 2, \ldots$.: τ is a stopping time with respect to this sequence if $\{\tau = n\}$ is completely determined by (at most) the total information known up to time n, $\{X_1, \ldots, X_n\}$.*

1.2 Examples

1. *(First passage/Hitting times)* Given a state i from the state space of $\mathbf{X}$, let

$$\tau = \min\{n \geq 0 : X_n = i\}.$$

   This is called the first passage time of the process into state i. Also called the *hitting time* of the process to state i. More generally we can let A be a collection of states such as $A = \{2, 3, 9\}$ or A=$\{2, 4, 6, 8, \ldots\}$, and then τ is the first passage time (hitting time) into the set A:

$$\tau = \min\{n \geq 0 : X_n \in A\}.$$

   An example is the gambler's ruin problem in which gambling stops when either $X_n = N$ or $X_n = 0$ whichever happens first; here $A = \{0, N\}$.

Proving that hitting times are stopping times is simple:

$\{\tau = 0\} = \{X_0 \in A\}$, hence only depends on X_0, and for $n \geq 1$,

$$\{\tau = n\} = \{X_0 \notin A, \ldots, X_{n-1} \notin A, X_n \in A\},$$

and thus only depends on $\{X_0, \ldots, X_n\}$ as is required. When $A = \{i\}$ this reduces to the hitting time to state i and

$$\{\tau = n\} = \{X_0 \neq i, \ldots, X_{n-1} \neq i, X_n = i\}.$$

In the gambler's ruin problem with $X_0 = i \in \{1, \ldots, N-1\}$, $X_n = i + \Delta_1 + \cdots + \Delta_n$ (simple random walk) until the set $A = \{0, N\}$ is hit. Thus $\tau = \min\{n \geq 0 : X_n \in A\}$ is a stopping time with respect to both $\{X_n\}$ and $\{\Delta_n\}$. For example, for $n \geq 2$,

$$\begin{aligned}\{\tau = n\} &= \{X_0 \notin A, \ldots, X_{n-1} \notin A, X_n \in A\}, \\ &= \{i + \Delta_1 \notin A, \ldots, i + \Delta_1 + \cdots + \Delta_{n-1} \notin A, i + \Delta_1 + \cdots + \Delta_n \in A\},\end{aligned}$$

and thus is compltely determined by both $\{X_0, \ldots, X_n\}$ and $\{\Delta_1, \ldots, \Delta_n\}$. The point here is that if we know the intial condition $X_0 = i$, then $\{X_n\}$ and $\{\Delta_n\}$ contain the same information.

2. *(Independent case)* Let $\mathbf{X} = \{X_n : n \geq 0\}$ be any stochastic process and suppose that τ is any random time that is independent of $\mathbf{X}$. Then τ is a stopping time. In this case, $\{\tau = n\}$ doesn't depend *at all* on $\mathbf{X}$ (past or future); it is independent of it. An example might be: Before you begin gambling you decide that you will stop gambling after the 10th gamble (regardless of all else). In this case $P(\tau = 10) = 1$. Another example: Every day after looking at the stock price, you flip a coin. You decide to sell the stock the first time that the coin lands heads. (I do not recommend doing this!) In this case τ is independent of the stock pricing and has a geometric distribution.

3. *(Example of a non-stopping time: Last exit time)* Consider the rat in the open maze problem in which the rat eventually reaches freedom (state 0) and never returns into the maze. Assume the rat starts off in cell 1; $X_0 = 1$. Let τ denote the last time that the rat visits cell 1 before leaving the maze:

$$\tau = \max\{n \geq 0 : X_n = 1\}.$$

Clearly we need to know the future to determine such a time. For example the event $\{\tau = 0\}$ tells us that in fact the rat never returned to state 1: $\{\tau = 0\} = \{X_0 = 1, X_1 \neq 1, X_2 \neq 1, X_3 \neq 1, \ldots\}$. Cleary this depends on all of the future, not just X_0. Thus this is not a stopping time.

In general a last exit time (the last time that a process hits a given state or set of states) is not a stopping time; inorder to know that the last visit has just occured, one must know the future.

1.3 Other formulations for stopping time

If τ is a stopping time with respect to $\{X_n\}$, then we can conclude that the event $\{\tau \leq n\}$ can only depend at most on $\{X_0, \ldots, X_n\}$: stopping by time n can only depend on the information up to time n. Formally we can prove this as follows: $\{\tau \leq n\}$ is the union of $n + 1$ events

$$\{\tau \leq n\} = \cup_{j=0}^{n} \{\tau = j\}.$$

By the definition of stopping time, each $\{\tau = j\}$, $j \leq n$, depends (at most) on $\{X_0, \dots, X_j\}$ which is contained in $\{X_0, \dots, X_n\}$. Thus the union is also contained in $\{X_0, \dots, X_n\}$.

Similarly we can handle a set like $\{\tau < n\}$ since we can re-write it as $\{\tau \leq n-1\}$; thus it is determined by $\{X_0, \dots, X_{n-1}\}$. Also, we can handle a set like $\{\tau > n\}$, since it is equivalent to $\overline{\{\tau \leq n\}}$, denoting the complement of the event $\{\tau \leq n\}$: since $\{\tau \leq n\}$ is determined by $\{X_0, \dots, X_n\}$, so is its complement. For example, if $\tau = \min\{n \geq 0 : X_n = i\}$, a hitting time, then $\{\tau > n\} = \{X_0 \neq i, X_1 \neq i, \dots, X_n \neq i\}$, and hence only depends on $\{X_0, \dots, X_n\}$.

1.4 Wald's Equation

We now consider the very special case of stopping times when $\{X_n : n \geq 1\}$ is an independent and identically distributed (i.i.d.) sequence with common mean $E(X)$. We are interested in the sum of the r.v.s. up to time τ,

$$\sum_{n=1}^{\tau} X_n = X_1 + \cdots + X_\tau.$$

Theorem 1.1 (Wald's Equation) *If τ is a stopping time with respect to an i.i.d. sequence $\{X_n : n \geq 1\}$, and if $E(\tau) < \infty$ and $E(|X|) < \infty$, then*

$$E\{\sum_{n=1}^{\tau} X_n\} = E(\tau)E(X).$$

This is simply a generalization of the fact that for any fixed integer n

$$E(X_1 + \cdots + X_n) = nE(X).$$

Wald's equation allows us to replace deterministic time n by the expected value of a random time τ when τ is a stopping time.

Proof :

$$\{\sum_{n=1}^{\tau} X_n\} = \{\sum_{n=1}^{\infty} X_n I\{\tau > n-1\}\},$$

where $I\{\tau > n-1\}$ denotes the indicator r.v. for the event $\{\tau > n-1\}$. By the definition of stopping time, $\{\tau > n-1\}$ can only depend (at most) on $\{X_1, \dots, X_{n-1}\}$ (Recall Section 1.3.) Since the sequence is assumed i.i.d., X_n is independent of $\{X_1, \dots, X_{n-1}\}$ so that X_n is independent of the event $\{\tau > n-1\}$ yielding $E\{X_n I\{\tau > n-1\}\} = E(X)P(\tau > n-1)$. Taking the expected value of the above infinite sum thus yields (after bringing the expectation inside the sum; that's allowed here since $E(\tau)$ and $E|X|$ are assumed finite)

$$\begin{aligned} E\{\sum_{n=1}^{\tau} X_n\} &= E(X)\sum_{n=1}^{\infty} P(\tau > n-1) \\ &= E(X)\sum_{n=0}^{\infty} P(\tau > n) \\ &= E(X)E(\tau), \end{aligned}$$

where the last equality is due to "integrating the tail" method for computing expected values of non-negative r.v.s. (Lecture Notes 1). ■

1.5 Applications of Wald's equation

1. Consider any i.i.d. sequence $\{X_n\}$ with $E(X) = 5.5$. Let τ be any random time that is independent of $\{X_n\}$ and has mean $E(\tau) = 10$. Then since it is independent it automatically is a stopping time, and Wald equation yields

$$E\{\sum_{n=1}^{\tau} X_n\} = E(\tau)E(X) = (10)(5.5) = 55.$$

2. Consider an i.i.d. sequence $\{X_n\}$ with a discrete distribution that is uniform over the integers $\{1, 2, \ldots, 10\}$; $P(X = i) = 1/10,\ 1 \le i \le 10$. Thus $E(X) = 5.5$. Imagine that these are bonuses (in some unit of money) that are given to you by your employer each year. Let $\tau = \min\{n \ge 1 : X_n = 6\}$, the first time that you receive a bonus of size 6.

   What is the expected total (cumulative) amount of bonus received up to time τ?

$$E\{\sum_{n=1}^{\tau} X_n\} = E(\tau)E(X) = 5.5E(\tau),$$

   from Wald's equation, if we can show that τ is a stopping time with finite mean.

   That τ is a stopping time follows since it is a first passage time: $\{\tau = 1\} = \{X_1 = 6\}$ and in general $\{\tau = n\} = \{X_1 \ne 6, \ldots, X_{n-1} \ne 6, X_n = 6\}$ only depends on $\{X_1, \ldots, X_n\}$.

   We need to calculate $E(\tau)$. Noting that $P(\tau = 1) = P(X_1 = 6) = 0.1$ and in general, from the i.i.d. assumption placed on $\{X_n\}$,

$$\begin{aligned} P(\tau = n) &= P(X_1 \ne 6, \ldots, X_{n-1} \ne 6, X_n = 6) \\ &= P(X_1 \ne 6) \cdots P(X_{n-1} \ne 6)P(X_n = 6) \\ &= (0.9)^{n-1}0.1,\ n \ge 1, \end{aligned}$$

   we conclude that τ has a geometric distribution with "success" probability $p = 0.1$, and hence $E(\tau) = 1/p = 10$. And our final answer is $E(\tau)E(X) = 55$.

   Note here that before time $\tau = n$ the random variables $X_1, \ldots, X_{n-1}$ no longer have the original uniform distribution; they are biased in that none of them takes on the value 6. So in fact they each have the conditional distribution $(X|X \ne 6)$ and thus an expected value different from 5.5. Moreover, the random variable at time $\tau = n$ has value 6; $X_n = 6$ and hence is not random at all. The point here is that even though all these random variables are biased, in the end, on average, Wald's equation let's us treat the sum as if they are not biased and are independent of τ as in Example 1 above.

   To see how interesting this is, note further that we would get the same answer 55 by using any of the stopping times $\tau = \min\{n \ge 1 : X_n = k\}$ for any $1 \le k \le 10$; nothing special about $k = 6$.

   This should indicate to you why Wald's equation is so important and useful.

3. (*Null recurrence of the simple symmetric random walk*)

   $X_n = \Delta_1 + \cdots + \Delta_n,\ X_0 = 0$ where $\{\Delta_n : n \ge 1\}$ is i.i.d. with $P(\Delta = \pm 1) = 0.5$, $E(\Delta) = 0$. We already know that this MC is recurrent (proved via the gambler's ruin problem and alternatively using Sterling's formula), but now we show that it is null recurrent. We do so by proving that $E(\tau_{1,1}) = \infty$. (Since the chain is irreducible, all states are null recurrent together or positive recurrent together; so if $E(\tau_{1,1}) = \infty$ then in fact $E(\tau_{j,j}) = \infty$ for all j.) In fact by symmetry, $\tau_{j,j}$ has the same distribution (hence mean) for all j.

Let $\tau_{i,j} = \min\{n \geq 1 : X_n = j | X_0 = i\}$. Note that each of these random times is in fact a stopping time with respect to the i.i.d. sequence $\{\Delta_n : n \geq 1\}$. By conditioning on the first step $\Delta_1 = \pm 1$,

$$\begin{aligned} E(\tau_{1,1}) &= (1 + E(\tau_{2,1}))1/2 + (1 + E(\tau_{0,1}))1/2 \\ &= 1 + 0.5E(\tau_{2,1}) + 0.5E(\tau_{0,1}). \end{aligned}$$

We will show that $E(\tau_{0,1}) = \infty$, thus proving the result.

Note that by definition, the chain at time $\tau = \tau_{0,1}$ has value 1;

$$1 = X_\tau = \sum_{n=1}^{\tau} \Delta_n.$$

Now we use Wald's equation with τ: for if in fact $E(\tau) < \infty$ then we conclude that

$$1 = E(X_\tau) = E\{\sum_{n=1}^{\tau} \Delta_n\} = E(\tau)E(\Delta) = 0,$$

yielding the contradiction $1 = 0$; thus $E(\tau) = \infty$.

1.6 Strong Markov property

Consider a Markov chain $\mathbf{X} = \{X_n : n \geq 0\}$ with transition matrix P. The Markov property can be stated as saying: *Given the state X_n at any time n (the present time), the future $\{X_{n+1}, X_{n+2}, \ldots\}$ is independent of the past $\{X_0, \ldots X_{n-1}\}$.*

If τ is a stopping time with respect to the Markov chain, then in fact, we get what is called the *Strong Markov Property*: *Given the state X_τ at time τ (the present), the future $\{X_{\tau+1}, X_{\tau+2}, \ldots\}$ is independent of the past $\{X_0, \ldots X_{\tau-1}\}$.*

The point is that we can replace a deterministic time n by a stopping time τ and retain the Markov property. It is a stronger statement than the Markov property and includes it as a special case: Choosing τ to be deterministic, $P(\tau = n) = 1$, yields the Markov property.

Theorem 1.2 *Any discrete-time Markov chain satisfies the Strong Markov Property.*

We actually have used this result already without saying so: Every time the rat in the maze retuned to cell 1 we said that "The chain starts over again and is independent of the past". Formally, we were using the strong Markov property with the stopping time

$$\tau = \tau_{1,1} = \min\{n \geq 1 : X_n = 1 | X_0 = 1\}.$$

The theorem holds (intuitively) because by definition of a stopping time, $\{\tau = n\}$ only depends on $\{X_0, \ldots, X_n\}$, the past and the present, and not on any of the future: Given the joint event $\{\tau = n, X_n = i\}$, the future $\{X_{n+1}, X_{n+2}, \ldots\}$ is still independent of the past.

Corollary 1.1 *As a consquence of the strong Markov property, we conclude that the chain from time τ onwards, $\{X_{\tau+n} : n \geq 0\}$, is itself the same Markov chain but with initial condition $X_0 = X_\tau$.*

The above makes perfect sense; for example in the context of the rat in the maze with $X_0 = 1$, we know that whenever the rat enters cell 2, the rats movement from then onwards is still the same Markov chain but with initial condition $X_0 = 2$.