

CS265/CME309: Randomized Algorithms and Probabilistic Analysis

Lecture #17: The Martingale Stopping Theorem, and Applications

Gregory Valiant*, updated by Mary Wootters

December 11, 2023

1 The Martingale Stopping Theorem

Given that a martingale, by definition, has the property that $\mathbf{E}[Z_t | Z_0, \dots, Z_{t-1}] = Z_{t-1}$, it follows inductively that $\mathbf{E}[Z_t] = \mathbf{E}[Z_0]$ for any fixed value of t . In this lecture, we consider what can happen if we consider $\mathbf{E}[Z_T]$ when T is a random variable. Before formally defining a stopping time, we begin with two examples that motivate the subtlety that can arise according to the properties of T .

Example 1. Consider the gambling game where we start with 0 dollars, and at every timestep, we either win or lose a dollar with probability $1/2$, independently of the outcome of previous timesteps. Let Z_t denote our net winnings/losses after t rounds of the above game. Consider the random variable T defined as the first time at which $|Z_t| = 10$, and the random variable T' define as the first time at which $Z_t = 10$. Both T and T' are perfectly valid random variables, because they are finite with probability 1, namely

$$\lim_{t \rightarrow \infty} \Pr[\text{exists } t' < t \text{ s.t. } |Z_{t'}| = 10] = 1, \text{ and } \lim_{t \rightarrow \infty} \Pr[\text{exists } t' < t \text{ s.t. } Z_{t'} = 10] = 1.$$

In the case of T , it holds that $\mathbf{E}[Z_T] = Z_0 = 0$. In the case of T' , however, by the definition of T' $\mathbf{E}[Z_{T'}] = 10 \neq Z_0$.

In the above example, what makes T and T' so different? More broadly, is there a set of properties that T could have that would guarantee that $\mathbf{E}[Z_T] = \mathbf{E}[Z_0]$? The *Martingale Stopping Theorem*, also sometimes referred to as the *Optional Stopping Theorem*, provides one set of sufficient such conditions.

Before stating the stopping theorem, we formalize the notion of a randomized “stopping time”. Informally, a stopping time T for a discrete time process $\{X_t\}$ is an integer-valued random variable with the property that the event that $T = i$ depends only on X_0, \dots, X_i and in particular is independent of X_{i+1}, \dots conditioned on X_0, \dots, X_i . Throughout we will also assume that the stopping

*©2019, Gregory Valiant. Not to be sold, published, or distributed without the authors' consent.

criteria occurs with probability 1. (If that doesn't hold, then T is a random variable that is undefined with positive probability. . . .)

Definition 2. Given a discrete time process $\{X_t\}$ (which can be a martingale, Markov chain, or anything else), a random variable, T , is a stopping time for $\{X_t\}$ if the event that $T = i$ is mutually independent of the random variables $\{X_j | X_0, \dots, X_i\}$, for all $j > i$.

The following example illustrates a few stopping times, and one random variable that is *not* a stopping time.

Example 3. Given the unbiased random walk where $Z_0 = 0$ and $Z_i = Z_{i-1} + 1$ with probability $1/2$ and $Z_i = Z_{i-1} - 1$ with probability $1/2$, consider the following quantities:

- $T_1 = \min\{t : Z_t = 10\}$,
- $T_2 = \min\{t : Z_t \notin [-5, 10]\}$,
- $T_3 = \max\{t : Z_t > 0\}$.

T_1 and T_2 are perfectly good stopping times. T_3 , on the other hand, is not a stopping time, because determining if $T_3 = i$ involves looking into the future.

1.1 Statement of the Theorem

The following theorem provides one set of sufficient conditions for $\mathbf{E}[Z_T] = \mathbf{E}[Z_0]$ for a stopping time T . Note that these are sufficient conditions, but they are not necessary—there exist more general formulations of this theorem.

Theorem 1 (Martingale Stopping Theorem). Letting $\{Z_t\}$ denote a martingale with respect to $\{X_t\}$, and T a stopping time for $\{X_t\}$ that occurs with probability 1, then $\mathbf{E}[Z_T] = \mathbf{E}[Z_0]$ if at least one of the following conditions hold:

1. If there exists a constant c such that for all i , $|Z_i| < c$.
2. If there exists a constant c such that with probability 1, $T < c$.
3. If $\mathbf{E}[T] < \infty$, and there exists a constant c such that for all i , $\mathbf{E}[|Z_{i+1} - Z_i| | X_0, \dots, X_i] < c$.

The proof of the above theorem is not too long, but involves a careful analysis of which limits exist, etc. We have not covered this sort of thing in any rigor in this class, and we'll leave the proof of this for a class on measure theory.

Returning to Example 1, it should now be clear the difference between T and T' . While they are both valid stopping times, $\mathbf{E}[T']$ is *not* bounded, nor is T' , and hence the above theorem does not apply to T' . T on the other hand satisfies the first condition, as $|Z_i| \leq 10$. (Technically, to make this statement actually true, we could define the martingale $\{Z'_t\}$ by $Z'_t = Z_t$ for all $t \leq T$, and $Z'_t = Z_T$ for $t > T$, in which case $\mathbf{E}[Z_T] = \mathbf{E}[Z'_T]$, and $|Z'_t| \leq 10$ for all t .)

2 Hitting Times of Random Walks

One useful application of the Martingale Stopping Theorem is in bounding the amount of time it takes for certain events to happen. We illustrate this with two natural examples.

Theorem 2. Consider the unbiased random walk $\{Z_t\}$ where $Z_0 = 0$ and $Z_t = Z_{t-1} \pm 1$, with each outcome chosen with probability $1/2$ independently of previous steps. Letting T denote the first time that Z_t reaches either $-a$ or $+b$, then

$$\Pr[Z_T = b] = \frac{a}{a+b} \quad \text{and} \quad \mathbf{E}[T] = ab.$$

Proof. For the first part of the theorem, note that the martingale stopping theorem applies to Z_T , as the sequence is contained in range $[-a, b]$, and hence the first condition is satisfied. Hence $\mathbf{E}[Z_T] = Z_0 = 0$. On the other hand $0 = \mathbf{E}[Z_T] = b \Pr[Z_T = b] + (-a)(1 - \Pr[Z_T = b])$. Solving for $\Pr[Z_T = b]$, we get $\Pr[Z_T = b] = \frac{a}{a+b}$. (This result has the nice interpretation that the probability we walk away having won, is proportionate to how much we are willing to lose before we quit :)

To prove the second part, establishing the expected time until the game ends, we will combine the above result with an analysis of a *new* martingale. Consider the sequence $\{Y_t\}$ defined for $t = 0, \dots, T$, as $Y_t = Z_t^2 - t$. First we establish that $\{Y_t\}$ is a martingale with respect to $\{Z_t\}$, and then we will apply the martingale stopping theorem to $\{Y_t\}$. To check that $\{Y_t\}$ is a martingale, consider

$$\mathbf{E}[Y_t | Z_0, \dots, Z_{t-1}] = \frac{1}{2}(Z_{t-1} + 1)^2 + \frac{1}{2}(Z_{t-1} - 1)^2 - t = Z_{t-1}^2 - (t-1) = Y_{t-1}.$$

Neither of the first two conditions of the martingale stopping theorem are satisfied by $\{Y_t\}$, since it need not be bounded, and T is not bounded. The third condition, however, holds because $|Y_{t+1} - Y_t| \leq 1 + 2|Z_t| = 1 + 2 \max(a, b)$, and $\mathbf{E}[T]$ is bounded, as can be seen by noting that, no matter the value of Z_t , there with probability at least $1/2^{b+a}$, the game will end after $b+a$ additional steps, hence $\mathbf{E}[T] \leq 2^{b+a}(b+a) < \infty$.

Applying the stopping theorem yields that $0 = Y_0 = \mathbf{E}[Y_T] = \mathbf{E}[Z_T^2] - \mathbf{E}[T]$. Now, we use the fact that we know everything there is to know about Z_T , namely we know that $Z_T = -a$ with probability $b/(a+b)$ and $Z_T = b$ with probability $a/(a+b)$. Hence

$$\mathbf{E}[T] = \mathbf{E}[Z_T^2] = a^2 \frac{b}{a+b} + b^2 \frac{a}{a+b} = ab \frac{a+b}{a+b} = ab.$$

□

The trickiest step in the proof of the above theorem was the construction of the martingale $Y_t = Z_t^2 - t$. While it is easy to verify that this is a martingale, the intuition for why an expression like this would be a martingale is not intuitive. The following theorem, characterizing the behavior of biased random walks, leverages an even more mysterious martingale. The statement of the following theorem is not the important part—the important part is that it is possible to leverage the martingale stopping theorem to *exactly* calculate properties of stochastic processes, such as the expected time until something happens, the probabilities of certain outcomes, etc.

Theorem 3. *Suppose you start with 0 dollars, and at every timestep you win a dollar with probability p , and lose a dollar with probability $1 - p$. Letting T denote the first time that you are either up by b dollars, or down by a dollars, and Z_T denote your net winnings/losses at this time,*

$$\Pr[Z_T = -a] = (1 - \Pr[Z_T = b]) = \frac{1 - c^b}{c^{-a} - c^b},$$

where $c = \frac{1}{p} - 1$. Additionally,

$$\mathbf{E}[T] = \frac{(-a)\frac{1-c^b}{c^{-a}-c^b} + b(1 - \frac{1-c^b}{c^{-a}-c^b})}{2p - 1}.$$

Proof. Let $\{X_t\}$ denote the sequence of outcomes of the game, and Z_t denote the net winnings/losses through time t . Since the game is biased, $\{Z_t\}$ is *not* a martingale. We will now construct a martingale $\{Y_t\}$ where $Y_t = c^{Z_t}$ for an appropriate constant c , defined as a function of p , and then apply the stopping theorem to $\{Y_t\}$.

To determine the appropriate constant c , we want $Y_{t-1} = \mathbf{E}[Y_t|Z_0, \dots, Z_{t-1}]$, so we will simply calculate the expression on the right side, and hope that we can find a value of c to satisfy this equation.

$$\mathbf{E}[Y_t|Z_0, \dots, Z_{t-1}] = pc^{Z_{t-1}+1} + (1-p)c^{Z_{t-1}-1} = c^{Z_{t-1}} \left(cp + \frac{1-p}{c} \right).$$

So, to calculate the value of c as a function of p for which $\{Y_t\}$ is a martingale, we simply solve the equation $1 = cp + \frac{1-p}{c}$, yielding that c is a solution to $c^2p - c + (1-p) = 0$. There are two solutions to this quadratic, as long as $p \in (0, 1)$. One solution is $c = 1$, in which case we get a useless martingale where $Y_t = 1$ for all t . The other solution is $c = \frac{1}{p} - 1$, which yields a non-trivial martingale that we can use. Henceforth, we define $c = \frac{1}{p} - 1$.

Letting T denote the first time Z_t reaches either $-a$ or $+b$, by the Martingale Stopping Theorem, $\mathbf{E}[Y_T] = Y_0 = c^0 = 1$. On the other hand, $\mathbf{E}[Y_T] = c^{-a} \Pr[Z_T = -a] + c^b \Pr[Z_T = b]$. Setting this equal to 1 and solving for $\Pr[Z_T = -a]$ yields

$$\Pr[Z_T = -a] = \frac{1 - c^b}{c^{-a} - c^b},$$

which simplifies to $c^b/(c^b + 1)$ when $a = b$.

To calculate $\mathbf{E}[T]$, we will use the above result, together with a new martingale: Let $Q_t = Z_t - \alpha t$. We will define an appropriate constant α such that $\{Q_t\}$ is a martingale with respect to $\{Z_t\}$. We want $\mathbf{E}[Q_t|Z_0, \dots, Z_{t-1}] = Q_{t-1}$, so since $\mathbf{E}[Q_t] = Q_{t-1} + p(1) + (1-p)(-1) + \alpha$, we can simply set $\alpha = 1 - 2p$ to obtain a martingale. Applying the martingale stopping theorem, we have that $\mathbf{E}[Q_T] = (-a) \Pr[Z_T = -a] + b \Pr[Z_T = b] + \alpha \mathbf{E}[T] = Q_0 = 0$. From above, we have already calculated $\Pr[Z_T = -a]$, and so we can just plug in that expression and solve, yielding

$$\mathbf{E}[T] = \frac{(-a) \Pr[Z_T = -a] + b \Pr[Z_T = b]}{-\alpha} = \frac{(-a)\frac{1-c^b}{c^{-a}-c^b} + b(1 - \frac{1-c^b}{c^{-a}-c^b})}{2p - 1}.$$

□




Figure 1: Illustration of $\mathbf{E}[T]$ for the case that $a = b = 20$, and $a = b = 100$ as the bias p ranges from 0 to 1, illustrating that the expected time rapidly changes from quadratic in a to linear, as the bias deviates from $1/2$.

While this expression for $\mathbf{E}[T]$ is exact, it is hard to parse. Figure 1 depicts this function in the case that that $a = b$, illustrating the fact that when $p = 1/2$, $\mathbf{E}[T] = a^2$, whereas as p approaches 1, or 0, this time approaches a . Unsurprisingly, even for p that is fairly close to $1/2$, the bias is strong enough that $\mathbf{E}[T]$ is essentially linear in a , as opposed to quadratic.

Note: At this point we are done with the material for the before-class minilectures. The stuff after this point is for reference after class.

3 The Ballot Counting Theorem

One famous application of the Martingale Stopping Theorem is the following surprisingly crisp theorem, capturing the process of tallying a set of votes.

Theorem 4. *Suppose we have a set of n votes, with each vote being for either candidate A or B . Let $N_A, N_B = n - N_A$ denote the respective number of votes for each of the candidates. Assuming $N_A > N_B$, if the votes are tallied in a random order the probability that candidate A is ahead at all timesteps during the vote count is $\frac{N_A - N_B}{N_A + N_B}$.*

There is a purely combinatorial proof of the above theorem that analyzes the fraction of the $n!$ orderings in which the votes are counted, which respect the constraint that the partial tallies all have A ahead of B . Instead of the combinatorial proof, we will instead give a proof via martingales.

Proof. Let X_t denote the difference between the number of votes for candidate A versus for B after t total votes have been tallied. Now consider the sequence $\{Z_t\}$ defined by $Z_t = \frac{X_{n-t}}{n-t}$ which counts “backwards” from the final count. In particular, $Z_0 = \frac{N_A - N_B}{n}$.

We claim that $\{Z_t\}$ is a martingale (with respect to itself). Before proving this statement, we first finish the proof of the theorem assuming that the Z_t ’s form a martingale.

We define the stopping time T to be the smallest t for which $Z_t = 0$ if such a t exists, and otherwise set $T = n - 1$. Because there are just a finite number of terms, the stopping theorem holds, hence $\mathbf{E}[Z_T] = \mathbf{E}[Z_0] = \frac{N_A - N_B}{n}$.

There are two cases:

- Suppose that candidate A is always ahead. Then $T = n - 1$, and we have

$$Z_T = \frac{X_1}{1} = 1,$$

since $X_1 = 1 - 0 = 1$. (Here we are using the fact that A was always ahead, so it must be that the the first vote counted was for A too).

- If A was not always ahead in the count, then $T < n - 1$ and by definition $Z_T = 0$.

Hence letting p denote $\Pr[A \text{ always ahead}]$, we have that

$$\mathbf{E}[Z_T] = 1 \cdot p + 0 \cdot (1 - p) = \frac{N_A - N_B}{n},$$

where the second equality is from the martingale stopping theorem. Solving for p yields the theorem.

To finish the proof, we simply need to establish the $\{Z_t\}$ is, in fact, a martingale. This is simply a calculation, leveraging our definition of Z_t . To compute $\mathbf{E}[Z_t | Z_0, \dots, Z_{t-1}]$, recall that $Z_{t-1}(n - t + 1) = X_{n-t+1}$ is the difference between the number of votes for A and B after $n - t + 1$ votes have been counted. Hence, of these votes, $\alpha = \frac{n-t+1+X_{n-t+1}}{2}$ were for A and $\beta = \frac{n-t+1-X_{n-t+1}}{2}$ were for B . Since the sequence $\{Z_t\}$ is tallying the votes “backwards”, with probability $\alpha/(\alpha + \beta)$, the $n - t$ th vote counted was for A and with the remaining probability, it was for B . Hence

$$\begin{aligned} \mathbf{E}[Z_t | Z_0, \dots, Z_{t-1}] &= \frac{X_{n-t+1} + 1}{n - t} \left(\frac{\frac{n-t+1+X_{n-t+1}}{2}}{n - t + 1} \right) + \frac{X_{n-t+1} - 1}{n - t} \left(\frac{\frac{n-t+1-X_{n-t+1}}{2}}{n - t + 1} \right). \\ &= \frac{X_{n-t+1}}{n - t} + \frac{X_{n-t+1}}{(n - t)(n - t + 1)} = \frac{X_{n-(t-1)}}{n - (t - 1)} = Z_{t-1}. \end{aligned}$$

□