The Martingale Stopping Theorem

Scott M. LaLonde

February 27, 2013

Abstract

We present a proof of the Martingale Stopping Theorem (also known as Doob's Optional Stopping Theorem). We begin with some preliminaries on measure-theoretic probability theory, which allows us to discuss the definition and basic properties of martingales, We then state some auxiliary results and use them to prove the main theorem.

1 Introduction

Recall that a **martingale** is (informally) a random process $X = \{X_n\}$ which models a player's fortune in a fair game. That is to say, his expected fortune at time n given the history up to this point is equal to his current fortune:

$$\mathbf{E}(X_n|X_1, \dots, X_{n-1}) = X_{n-1}.$$

This in turn implies that for all n,

$$\mathbf{E}(X_n) = \mathbf{E}(X_{n-1}) = \cdots = \mathbf{E}(X_1) = \mathbf{E}(X_0),$$

so the player's expected fortune at any time is equal to his starting expected fortune. It is natural to ask whether the game remains fair when stopped at a randomly chosen time. Loosely speaking, if T is a random stopping time and X_T denotes the game stopped at this time, do we have

$$\mathbf{E}(X_T) = \mathbf{E}(X_0)$$

as well? In general the answer is no, as Doyle and Snell point out. They envision a situation where the player is allowed to go into debt by any amount and to play for an arbitrarily long time. In such a situation, the player will inevitably come out ahead. There are conditions which will guarantee fairness, and Doyle and Snell [2] give them in the following theorem, which is phrased in the context of gambling.

Theorem (Martingale Stopping Theorem)**.** *A fair game that is stopped at a random time will remain fair to the end of the game if it is assumed that:*

(a) There is a finite amount of money in the world.

(b) A player must stop if he wins all of this money or goes into debt by this amount.

Our goal is to develop a more formal statement of this theorem, called Doob's Optional-Stopping Theorem, and then to prove it. We will start with some general background material on probability theory, provide formal definitions of martingales and stopping times, and finally state and prove the theorem. It should be noted that our exposition will largely be based on that of Williams [4], though a nice overview of martingales and various results about them can be found in Doob [1].

2 Preliminaries

Modern approaches to probability theory make much use of measure theory. Since the proof of Doob's theorem will rely heavily on some sort of integral convergence theorem (namely the Dominated Convergence Theorem), we need to introduce some background that places probability theory within the realm of measure theory.

In modern probability theory the model for a random experiment is called a *probability space*. This is a triple $(\Omega, \Sigma, \mathbf{P})$, where

- Ω is a set, called the *sample space*.
- Σ is a σ-algebra of subsets of Ω.
- $\mathbf{P}$ is a probability measure on (Ω, Σ), i.e. every set in Σ is measurable and

$$\mathbf{P}(\Omega) = 1.$$

The notion of a probability space generalizes ideas from discrete probability. We have already mentioned that Ω is the *sample space* of an experiment. The σ-algebra Σ represents the set of possible outcomes, or the *events* to which one can assign a probability. The measure $\mathbf{P}$ gives the probability that an outcome occurs.

Of course in discrete probability one is usually interested in *random variables*, which are real-valued functions on the sample space. For us, a random variable will be a function $X : \Omega \to \mathbf{R}$ which is measurable with respect to Σ. The *expected value* of a random variable X is its integral with respect to the measure $\mathbf{P}$:

$$\mathbf{E}(X) = \int_{\Omega} X(\omega)\, d\mathbf{P}(\omega),$$

and we will say that a random variable X is *integrable* if $\mathbf{E}(|X|) < \infty$. Finally, we will need to make reference to the *conditional expectation* of a random variable: given a sub-σ-algebra $\mathcal{A}$ of Σ, the conditional expectation $\mathbf{E}(X|\mathcal{A})$ is a random variable which satisfies certain conditions related to X and $\mathcal{A}$. The proper definition is quite complicated, so one should simply think of $\mathbf{E}(X|\mathcal{A})$ as the expectation of X given that the events contained in $\mathcal{A}$ have occurred. This description is not completely accurate, but it should help the reader's understanding of the uses of conditional expectation in the sequel.

3 Martingales and Stopping Times

Now that we have the appropriate background material out of the way, we can formally define a martingale. Fix a probability space $(\Omega, \Sigma, \mathbf{P})$, and let $X = \{X_n\}_{n=0}^{\infty}$ be a sequence of random variables on Ω.

Definition 1. *A **filtration** on $(\Omega, \Sigma, \mathbf{P})$ is an increasing sequence $\mathcal{F} = \{F_n\}_{n=1}^{\infty}$*

$$F_1 \subset F_2 \subset F_3 \subset \cdots \subset \Sigma$$

*of sub-σ-algebras of Σ. The sequence $\{X_n\}_{n=1}^{\infty}$ is said to be **adapted** to $\mathcal{F}$ if X_n is F_n-measurable for each n.*

Remark 2. This definition may seem abstract, but it helps to keep the following idea in mind. The σ-algebra F_n represents the information available to us at time n in a random process, or the events that we can detect at time n. That the sequence is increasing represents the fact that we gain information as the process goes on.

This idea can perhaps be made even more clear by pointing out that a common choice for $\mathcal{F}$ is the *natural filtration* (or *minimal filtration*):

$$F_n = \sigma(X_1, \ldots, X_n).$$

In this case, F_n is the smallest σ-algebra on Ω making the random variables $X_1, \ldots, X_n$ measurable. The information available at time n is precisely that generated by the X_i for $1 \leq i \leq n$. Of course more information could be available, which would correspond to a different choice of filtration $\mathcal{F}$.

Filtrations are important because they provide a concise way of defining a martingale. With this in mind, let $\mathcal{F} = \{F_n\}$ be a fixed filtration on $(\Omega, \Sigma, \mathbf{P})$.

Definition 3. *A random process $X = \{X_n\}$ is called a **martingale** relative to $\mathcal{F}$ if*

(a) X is adapted to $\mathcal{F}$,

(b) $\mathbf{E}(|X_n|) < \infty$ for all n, and

(c) $\mathbf{E}(X_{n+1}|F_n) = X_n$ almost surely.

As we have already discussed, we are interested in what happens when one stops a martingale at a random time. To do this, we need a formal way of talking about a rule for stopping a random process which does not depend on the future. This leads to the following definition of a *stopping time*.

Definition 4. *A map $T : \Omega \to \{1, 2, \ldots, \infty\}$ is called a **stopping time** if*

$$\{T = n\} = \{\omega \in \Omega : T(\omega) = n\} \in F_n \tag{1}$$

*for all $n \leq \infty$. We will say that T is **almost surely finite** if $\mathbf{P}(\{T = \infty\}) = 0$.*

Remark 5. Intuitively, T is a random variable taking positive integer values (and possibly ∞) which gives a rule for stopping a random process. Condition (1) says that the decision whether to stop or not at time n depends only on the information available to us at time n (i.e. the history up to and including time n). No knowledge of the future is required, since such a rule would surely result in an unfair game.

Let $X = \{X_n\}$ be a random process, and let T be a stopping time. For any positive integer n and any $\omega \in \Omega$, we define

$$T \wedge n(\omega) = \min\{T(\omega), n\}.$$

With this notation, we can define a *stopped process.*

Definition 6. *The **stopped process** $X^T = \{X_n^T\}$ is given by*

$$X_n^T(\omega) = X_{T \wedge n(\omega)}(\omega).$$

A useful result that we will need for the proof of Doob's theorem (but that we will not prove) says that X^T inherits certain desirable properties from X.

Proposition 7. *If $X = \{X_n\}$ is a martingale, then the stopped process $X^T = \{X_{T \wedge n}\}$ is also a martingale. In particular, for all n we have*

$$\mathbf{E}(X_{T \wedge n}) = \mathbf{E}(X_0).$$

This is part (ii) of [4, Theorem 10.9], and an outline of the proof can be found there. The proof is not difficult, but the details are not particularly enlightening from our current perspective.

4 Doob's Optional-Stopping Theorem

We now have all the pieces in place to state and prove our main theorem. First we need to formalize what it means to "stop a process at a random time." Suppose we have a martingale $X = \{X_n\}$ and a stopping time T. Assume that T is almost surely finite. Then we can define a random variable $X_T : \Omega \to \mathbf{R}$ by

$$X_T(\omega) = X_{T(\omega)}(\omega),$$

at least for ω outside some set of probability 0. (To make X_T everywhere-defined, we could set it equal to 0 on this null set.) Intuitively, $\mathbf{E}(X_T)$ represents the player's expected fortune when stopping at a random time. If we are to show that we still have a fair game, we will need to check that

$$\mathbf{E}(X_T) = \mathbf{E}(X_0). \tag{2}$$

Note that $T \wedge n$ converges to T pointwise almost surely as $n \to \infty$, so we have $X_{T\wedge n} \to X_T$ almost surely. Moreover, we know that $\mathbf{E}(X_{T\wedge n}) = \mathbf{E}(X_0)$ for all n. It would be nice if we could conclude that $\mathbf{E}(X_{T\wedge n}) \to \mathbf{E}(X_T)$, since we would then have (2). This amounts to showing that

$$\int_\Omega X_{T\wedge n}(\omega)\, d\mathbf{P}(\omega) \to \int_\Omega X_T(\omega)\, d\mathbf{P}(\omega)$$

as $n \to \infty$. Convergence of this sort is by no means guaranteed. We need hypotheses that will allow us to invoke convergence theorems from measure theory, with the Dominated Convergence Theorem being the likely candidate.

We will prove the version of Doob's theorem given in [4, Theorem 10.10], which is essentially the same as the formal statement given in class. The proof of part (b) will differ slightly from Williams' proof, however. In the process we will obtain direct analogues of the Martingale Stopping Theorem from [2]. In this regard, requirement that there is only "a finite amount of money in the world" can be encoded by assuming that the random variables X_n are uniformly bounded; this is condition (b) below. Similarly, the requirement that the player stop after a finite amount of time is obtained by requiring that T be almost surely bounded, which is condition (a). We also show that there is a third condition under which the theorem holds; this condition is essentially limit on the size of a bet at any given time.

Theorem 8 (Doob's Optional-Stopping Theorem)**.** *Let* $(\Omega, \Sigma, \mathbf{P})$ *be a probability space,* $\mathcal{F} = \{F_n\}$ *a filtration on* Ω*, and* $X = \{X_n\}$ *a martingale with respect to* $\mathcal{F}$*. Let* T *be a stopping time. Suppose that any one of the following conditions holds:*

(a) There is a positive integer N *such that* $T(\omega) \le N$ *for all* $\omega \in \Omega$*.*

(b) There is a positive real number K *such that*

$$|X_n(\omega)| < K$$

for all n *and all* $\omega \in \Omega$*, and* T *is almost surely finite.*

(c) $\mathbf{E}(T) < \infty$*, and there is a positive real number* K *such that*

$$|X_n(\omega) - X_{n-1}(\omega)| < K$$

for all n *and all* $\omega \in \Omega$*.*

Then X_T *is integrable, and*

$$\mathbf{E}(X_T) = \mathbf{E}(X_0).$$

Proof. Note that in all three cases T is a.s. finite. By our previous discussion, this implies that X_T is a.s.-defined, and we have $X_{T\wedge n} \to X_T$ almost surely. Furthermore, we know that $X_{T\wedge n}$ is integrable for all n, and that $\mathbf{E}(X_{T\wedge n}) = \mathbf{E}(X_0)$.

Suppose that (a) holds. Then for $n \geq N$, we have $T(\omega) \wedge n = T(\omega)$ for all $\omega \in \Omega$. Hence $X_{T\wedge n} = X_T$ for $n \geq N$, and it follows that X_T is integrable with

$$\mathbf{E}(X_T) = \mathbf{E}(X_{T\wedge N}) = \mathbf{E}(X_0).$$

Now suppose that (b) holds. Then the boundedness condition on the X_n implies that

$$|X_{T\wedge n}(\omega)| < K$$

for all n and all $\omega \in \Omega$. Also, it is fairly easy to check that

$$X_{T\wedge n}(\omega) = X_0(\omega) + \sum_{k=1}^{T\wedge n(\omega)} X_k(\omega) - X_{k-1}(\omega)$$

for all ω, so if (c) holds we have

$$|X_{T\wedge n}(\omega)| \leq |X_0(\omega)| + \sum_{k=1}^{T\wedge n(\omega)} |X_k(\omega) - X_{k-1}(\omega)| \leq |X_0(\omega)| + KT(\omega).$$

Certainly X_0 is integrable, and we have $\mathbf{E}(KT) = K\mathbf{E}(T) < \infty$ by assumption. Therefore, in either case (b) or (c) we have bounded $|X_{T\wedge n}|$ by an integrable random variable, so the Dominated Convergence Theorem applies. It follows that X_T is integrable, and

$$\lim_{n\to\infty} \int_\Omega X_{T\wedge n}(\omega)\, d\mathbf{P}(\omega) = \int_\Omega X_T(\omega)\, d\mathbf{P}(\omega).$$

Equivalently,

$$\lim_{n\to\infty} \mathbf{E}(X_{T\wedge n}) = \mathbf{E}(X_T).$$

But $\mathbf{E}(X_{T\wedge n}) = \mathbf{E}(X_0)$ for all n, so we have $\mathbf{E}(X_T) = \mathbf{E}(X_0)$, as desired. □

References

[1] J.L. Doob, *What is a martingale?*, Amer. Math. Monthly **78** (1971), no. 5, 451–463.

[2] Peter G. Doyle and J. Laurie Snell, *Random walks and electrical networks*, Carus Mathematical Monographs, Mathematical Association of America, Washington, D.C., 1984.

[3] Efe A. Ok, *Probability theory with economic applications*, https://files.nyu.edu/eo1/public/books.html.

[4] David Williams, *Probability with martingales*, Cambridge University Press, Cambridge, 1991.