Optional Stopping Theorem Research

Steven Hwang

October 2020

Introduction

The purpose of this paper is to introduce and prove the Optional Stopping Theorem (also known as Doob's Optional Stopping Theorem, the Martingale Stopping Theorem, or the Optional Sampling Theorem).

Informal Statement of Optional Stopping Theorem

Wikipedia's summarizes the Optional Stopping Theorem as follows: "Under certain conditions, the expected value of a martingale at a stopping time is equal to its initial expected value."

In other words, if $X_0, X_1, ...$ is a martingale, and X_T is a stopping time, the Optional Stopping Theorem presents the conditions under which $\mathbf{E}(X_T) = \mathbf{E}(X_0)$. Informally, these conditions are:

- There is a finite amount of money in the world.
- A player must stop if he wins all of this money or goes into debt by this amount.

Background

In order to work up to a formal statement of this theorem, let's begin with a bit of review. We will be working in the probability space $(\Omega, \Sigma, \mathbf{P})$, defined as follows.

sample space: a set Ω representing all possible outcomes

event space: a σ-algebra of Ω, i.e. a set Σ consisting of all possible subsets of Ω, with each subset of outcomes being known as an event

probability measure: a function $\mathbf{P} : \Sigma \rightarrow [0, 1]$ assigning every event a real-valued probability

We'll also be working with random variables and expected values, so we'll recall those definitions as well.

random variable: a function $X : \Omega \to \mathbb{R}$ assigning every outcome an associated real value

expected value: given a random variable X, its expected value $\mathbf{E}(X)$ is the "average" output of X across all outcomes in Ω weighted by probability,

$$\mathbf{E}(X) = \sum_{\omega \in \Omega} X(\omega)\mathbf{P}(\omega)$$

We then move on to some new concepts that will be important in understanding the Optional Stopping Theorem. Most notably are the notions of filtrations and martingales.

martingale: a sequence of random variables such that each successive variable equals the expected value of all of the previous variables; i.e., a sequence of random variables $X_1, X_2, \ldots$ such that

$$\mathbf{E}(X_{n+1} \mid X_1, \ldots, X_n) = X_n$$

which also implies that

$$\mathbf{E}(X_{n+1}) = \mathbf{E}(X_n) = \ldots = \mathbf{E}(X_1) = \mathbf{E}(X_0)$$

filtration: a filtration $\mathcal{F} = \{F_n\}$ is an increasing sequence of subsets of Σ, $F_1 \subset F_2 \subset \ldots \subset \Sigma$, where each subset of events F_n represents the information available (events known) at time step n in a random process (sequence of random variables), and the fact that $\{F_n\}$ is an increasing sequence represents how information can only be gained in successive time steps

We can redefine martingales by rewriting the condition that the expected value of the next random variable equals the previous random variable using filtrations.

martingale relative to a filtration: a random process $\{X_n\}$ is called a martingale relative to filtration $\mathcal{F}$ if it satisfies the following conditions:

- X is adapted to $\mathcal{F}$; that is X_n is F_n-measurable for all n; in other words, for every Borel subset $A \subset \mathbb{R}$, the preimage of A, $X^{-1}(A) = \{\omega \in \Omega : X(\omega) \in A\}$, which is a subset of Ω and thus an element of Σ, is also an element of F_n, $X^{-1}(A) \in F_n$
- $\mathbf{E}(|X|) < \infty$ for all n
- $\mathbf{E}(X_{n+1}|F_n) = X_n$ almost surely; that is, the probability that $\mathbf{E}(X_{n+1}|F_n) = X_n$ is 1

Lastly, we will need to recall the definition of stopping times, and look at some associated notions.

stopping time: given filtration $\mathcal{F}$, a random variable whose range is restricted to the nonnegative integers $T : \Omega \to \mathbb{N}$ is called a stopping time if, for all time steps $n \leq \infty$, T's preimage of n is an element of F_n, i.e.

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

so that T represents a rule determining when to stop random process $\{X_n\} = X_0, X_1, X_2, \ldots$, and the condition for T to be a stopping time means that T only determines whether or not to stop at time n, $\{T = n\}$, based on the information available at time n, F_n

almost surely finite: stopping time T is almost surely finite if $\mathbf{P}(T = \infty) = 0$.

stopped process: given a random process $\{X_n\}$ and a stopping time T, the stopped process $X^T = \{X_n^T\}$ is a sequence of random variables $X_n^T : \Omega \to \mathbb{R}$ given by:

$$X_n^T(\omega) = X_{\min(T(\omega),n)}(\omega)$$

so that stopped process X^T must repeat the same random variable $X_{T(\omega)}(\omega)$ after reaching a certain time $T(\omega) \leq n$, as dictated by the stopping time function T

stopping time of a random process: given a martingale $X = \{X_n\}$ and a stopping time T, the stopping time of random process $\{X_n\}$ is a random variable $X_T : \Omega \to \mathbb{R}$ defined by

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

which is also the value that the stopped process X^T settles on

With all of these ideas in mind, we are ready to formalize the Optional Stopping Theorem.

Formal Statement of Optional Stopping Theorem

Let $(\Omega, \Sigma, \mathbf{P})$ be a probability space, $\mathbb{F} = \{F_n\}$ a filtration on Ω, and $X = \{X_n\}$ a martingale with respect to $\mathbb{F}$. Let T be a stopping time. Then

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

if any of the following conditions hold:

- There exists a positive integer N such that $T(\omega) < N$ for all $\omega \in \Omega$. (Intuitively, the random process always stops after a certain number of steps.)

- There exists a positive real K such that $|X_n(\omega)| < K$ for all $n \in \mathbb{N}$ and all $\omega \in \Omega$, and T is almost surely finite. (Intuitively, the random process is bounded.)
- $\mathbf{E}(T) < \infty$, and there exists a positive real K such that $|X_n(\omega) - X_{n-1}(\omega)| < K$ for all $n \in \mathbb{N}$ and all $\omega \in \Omega$. (Intuitively, the random variables are close together.)

Proof

Define operation $\wedge$ as follows: given stopping time T and index n, random variable $T \wedge n : \Omega \to \mathbb{N}$ is defined by $T \wedge n(\omega) = \min(T(\omega), n)$ (recall that $T(\omega) \in \mathbb{N}$).

Because random process $\{X_n\}$ is a martingale, the expected value of any random variable in the process equals the expected value of the first variable; and as $T \wedge n$ has outputs in the naturals, we know that

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

for all $n \in \mathbb{N}$ and all $\omega \in \Omega$. We also know that $\lim_{n\to\infty} X_{T\wedge n} = X_T$ almost surely, but in order to show that $\mathbf{E}(X_T) = \mathbf{E}(X_0)$, it remains to show that

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

Breaking this down using the definitions of expected values, we need to show that

$$\lim_{n\to\infty} \sum_{\omega\in\Omega} X_{T\wedge n}(\omega)\mathbf{P}(\omega) = \sum_{\omega\in\Omega} X_T(\omega)\mathbf{P}(\omega)$$

However, the sample spaces Ω is usually infinite and continuous, so it's more accurate to represent the expected values using integrals rather than sums. Thus, we want to show that

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

In other words, showing that the expected value of a martingale at a stopping time equals its initial expected value has reduced down to showing that this particular integral converges. While such a convergence of the integral is not guaranteed in general, each of the three conditions given in the statement of the theorem provides a sufficient basis to show the convergence.

From Condition 1, it doesn't require too much machinery to show that $\mathbf{E}(X_T) = \mathbf{E}(X_0)$, so the line of logic will be presented here. Assuming that there exists N such that $T(\omega) < N$ for all $w \in \Omega$, then for all $n \geq N$,

$$T \wedge n(\omega) = \min(T(\omega), n) = T(\omega) \ (n \geq N)$$

because $T(\omega) < N \leq n$. Then, because the functions $T \wedge n$ and T are identical for $n \geq N$, it follows that for all $\omega \in \Omega$, the random variables which $T \wedge n(\omega)$ and $T(\omega)$ are indices of should also be identical:

$$X_{T \wedge n(\omega)} = X_{T(\omega)} \ (n \geq N)$$

And if the random variables are identical, their expected values should also be equal:

$$\mathbf{E}(X_{T \wedge n(\omega)}) = \mathbf{E}(X_{T(\omega)}) \ (n \geq N)$$

As noted previously, because $\{X_n\}$ is a martingale,

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

In addition, the definition of the stopping time of a random process was that $X_T(\omega) = X_{T(\omega)}(\omega)$. Hence, substituting and setting $n = N$, we see that

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

Conditions 2 and 3 take the aforementioned integral approach, showing that the integral converges after assuming either of the two statements. However, these proofs require heavier machinery in order to show that $\mathbf{E}(X_T) = \mathbf{E}(X_0)$, including key results from measure theory, such as the Dominated Convergence Theorem, which are not within the scope of this paper. Thus, these two parts are omitted, but the more complete lines of reasoning can be found in first resource listed below.

Resources

- https://math.dartmouth.edu/~pw/math100w13/lalonde.pdf
- https://en.wikipedia.org/wiki/Optional_stopping_theorem
- https://en.wikipedia.org/wiki/Martingale_(probability_theory)
- https://en.wikipedia.org/wiki/Stopping_time
- https://en.wikipedia.org/wiki/Fundamental_theorem_of_asset_pricing
- https://en.wikipedia.org/wiki/Adapted_process
- https://en.wikipedia.org/wiki/Stopped_process
- https://en.wikipedia.org/wiki/Almost_surely
- https://en.wikipedia.org/wiki/Measurable_function
- https://math.stackexchange.com/questions/508790/what-does-it-mean-by-mathcalf-measurab