OPTIMAL STOPPING THEOREM & APPLICATIONS

JULIA RYCHLIK

ABSTRACT. This paper offers an overview of the Optimal Stopping Theorem and the conditional probability that is necessary to prove and understand it. We first give an overview of the measure theoretic definition of conditional probability as well as its important properties. Then martingales are defined and their important properties noted. Then, stopping times are explained. The proof of the optimal stopping theorem starts with the finite horizon case, proved using the optional stopping theorem. Then two examples of applications are detailed. Finally, the infinite case is treated and proved, with an example given as well.

CONTENTS

1. INTRODUCTION

This paper introduces the optimal stopping theorem, the intuition behind it, its proof, and some of its applications. An understanding of basic measure theoretic probability including random variables, distributions, expectation, and independence is assumed, but measure theoretic conditional expectation is treated thoroughly in the paper.

The optimal stopping theorem comes into play in situations where we are observing a process occurring over times $\{1, . . . N\}$ (finite horizon) or $\{1, 2, . . .\}$ (general), and we receive a reward (or incur a cost) whose distribution is modeled by a random variable X_n that depends on the time n that we stopped at. The optimal stopping theorem helps us find a stopping rule that can tell us when to stop the process in order to maximize the expected reward (or minimize the expected cost). Critically, in order to be useful in applications, the stopping rule that comes from this theorem only relies on information occurring before the stop, mirroring the real life observation of processes.

Date: SEPTEMBER 3, 2025.

Section 2 gives a general definition of the measure theoretic definition of conditional probability, and details its important properties that are necessary to understand or prove the optimal stopping theorem. The main reason that conditional expectation is necessary, is because it allows us to define and work with martingales in section 3.

Martingales are sequences of random variables such that the best approximation of a random variable later on in the sequence conditioned on the information known up to a time n in the sequence is just the n^{th} random variable in the sequence. This concept allows us to mathematically represent the situations that the optimal stopping theorem applies to where we want to extrapolate information about the future based on what we know up to the present moment.

In section 4, we define and examine relevant properties of stopping times. Stopping times are the object we seek to solve for in optimal stopping problems, as they are defined as random variables T taking values in $\{1, 2, . . ., \}$ such that for some increasing sequence of σ-algebras $\{\mathcal{F}_n\}_{n=1}^{\infty}$, $\{T = n\}$ is $\mathcal{F}_n$ measurable, meaning that if the variable tells you to stop at time n, it only needed information up to time n which is exactly the situation of the optimal stopping theorem.

Section 5 lays out the proof of the optimal stopping theorem for finite horizon problems using all the details so far. This proof requires the optional stopping theorem which is also proved. This theorem has applications and implications of its own, although they are not the subject of this paper. Additionally, section 5 includes two examples of problems that can be solved by the finite horizon optimal stopping theorem.

The paper ends in section 6 with a treatment of the general optimal stopping theorem that can be applied to non finite horizon problems. This new version of the theorem is proved in depth. Then an outline of an example and its solution is provided.

The following is probability notation that is used throughout the paper:

- $(\Omega, \mathcal{F}, \mathbb{P})$ for a probability space, its sigma algebra, and the measure defined on it.
- $\mathbb{E}(X; A) = \mathbb{E}(X \cdot \chi_A)$
- $\mathbb{P}(A) = \mathbb{E}(\chi_A)$
- $\{X \in A\} = \{\omega \mid X(\omega) \in A\}$
- For random variables X and Y, $\mathbb{E}(Y|X) = \mathbb{E}(Y|\sigma(X))$

2. Conditional expectation

In discrete probability, the conditional probability of an event A given an event B of nonzero probability is simple:

$$\mathbb{P}(A|B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}.$$

However, in measure theoretic probability, we are often dealing with events of probability zero, so we must we must generalize the idea of conditional expectation beyond this simple definition. This is done by conditioning a random variable on a σ-algebra rather than on a single event. The outcome of this generalized form of conditional probability is a random variable rather than a number. We now define such a conception of conditional probability.

Definition 2.1. For a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and a real-valued random variable X defined on this probability space such that $\mathbb{E}|X|$ is finite, the conditional expectation of X with respect to the σ-algebra $\mathcal{G}$ is a $\mathcal{G}$ measurable random variable Y such that for every $A \in \mathcal{G}$,

$$\mathbb{E}(X; A) = \mathbb{E}(Y; A)$$

The following proposition demonstrates why this definition makes sense:

Proposition 2.2. *For X defined as above and for a σ-algebra $\mathcal{G}$ there exists a Y such that*

$$Y = \mathbb{E}(X|\mathcal{G}).$$

Moreover, this Y is unique up to sets of measure zero.

The full proof of this statement can be found in Chatterjee [1] on pages 93-95.

Intuitively, we can interpret this definition of conditional expectation as the closest approximation of X that is $\mathcal{G}$-measurable, i.e. the closest approximation of X when we are restricted to the information contained in the σ-algebra $\mathcal{G}$. This becomes helpful for the optimal stopping theorem where we are trying to decide when to stop based on the information that has occurred up to that point in time; conditional expectation is used to condition on the information that is available at the moment that one must decide whether to stop.

The following proposition gives an interesting relationship between traditional conditional expectation and the measure theoretic definition.

Proposition 2.3. *Let A and B be any events such that $\mathbb{P}(B) > 0$. Define $\mathcal{G} = \{B, B^C, \Omega, \emptyset\}$. Then*

$$\mathbb{P}(A|\mathcal{G}) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} \cdot \chi_B + \frac{\mathbb{P}(A \cap B^C)}{\mathbb{P}(B^C)} \cdot \chi_{B^C}$$

almost surely.

Proof. Since we have almost sure uniqueness of conditional expectation, we just need to show that for

$$Y := \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} \cdot \chi_B + \frac{\mathbb{P}(A \cap B^C)}{\mathbb{P}(B^C)} \cdot \chi_{B^C}$$

we have

$$\mathbb{E}(Y; B) = \mathbb{E}(\chi_A; B) \text{ and } \mathbb{E}(Y; B^C) = \mathbb{E}(X; B^C).$$

Starting with B, we have

$$\begin{aligned}
\mathbb{E}(Y; B) &= \mathbb{E}\left(\chi_B \left(\frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} \cdot \chi_B + \frac{\mathbb{P}(A \cap B^C)}{\mathbb{P}(B^C)} \cdot \chi_{B^C}\right)\right) \\
&= \mathbb{E}\left(\frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} \cdot \chi_B\right) \\
&= \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} \cdot \mathbb{E}(\chi_B) \\
&= \mathbb{P}(A \cap B) \\
&= \mathbb{E}(\chi_A \chi_B) \\
&= \mathbb{P}(A; B).
\end{aligned}$$

Replacing B with B^C we see that it holds for B^C.. □

We now list some important properties about conditional expectation without proof. All proofs can be found in Chatterjee [1] on pages 96-97. In all of the following statements, X is a real-valued random variable defined on $(\Omega, \mathcal{F}, \mathbb{P})$ with $\mathbb{E}(X)$ finite.

Proposition 2.4. *Conditional expectation is linear.*

Proposition 2.5. *Nonnegativity is preserved by conditional expectation.*

Corollary 2.6. *Conditional expectation is monotone i.e. if $X \leq Y$ then for any σ-algebra $\mathcal{G}$,*

$$\mathbb{E}(X|\mathcal{G}) \leq \mathbb{E}(Y|\mathcal{G})$$

Proposition 2.7. *If X is independent of a σ-algebra then*

$$\mathbb{E}(X|\mathcal{G}) = \mathbb{E}(X)$$

almost surely and if X is $\mathcal{G}$ measurable then

$$\mathbb{E}(X|\mathcal{G}) = X$$

almost surely.

Proposition 2.8. *If for a σ-algebra $\mathcal{G}$, Y is $\mathcal{G}$ measurable and X and XY are integrable then*

$$\mathbb{E}(XY|\mathcal{G}) = Y\mathbb{E}(X|\mathcal{G})$$

almost surely.

Proposition 2.9. *The tower property states that for any σ-algebras $\mathcal{G}' \subset \mathcal{G}$,*

$$\mathbb{E}(X|\mathcal{G}') = \mathbb{E}(\mathbb{E}(X|\mathcal{G})|\mathcal{G}')$$

All of these properties will be used to prove the optimal stopping theorem.

3. Martingales, submartingales, and supermartingales

Definition 3.1. For a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, a filtration is a sequence of σ-algebras $\{\mathcal{F}_n\}_{n=0}^{\infty}$ such that

$$\mathcal{F}_0 \subset \mathcal{F}_1 \subset \mathcal{F}_2, \subset \ldots$$

and each $\mathcal{F}_n \subset \mathcal{F}$. A sequence of random variables $\{X_n\}_{n=0}^{\infty}$ is adapted to a filtration if for each n, X_n is $\mathcal{F}_n$ measurable. A martingale is an adapted sequence of random variables $\{X_n\}_{n=0}^{\infty}$ such that for each n,

$$\mathbb{E}(X_{n+1}|\mathcal{F}_n) = X_n$$

almost surely.

Based on the intuitive idea of conditional expectation defined in section 2, a martingale is a sequence of random variables such that, at any given time n, the best approximation of the next random variable in the sequence, with only the information provided by the random variables up to time n, is just the n^{th} variable in the sequence.

Below is a simple yet important property of martingales, that is a result of the tower property of conditional expectation:

Proposition 3.2. *By the tower property of conditional expectation, for a martingale $\{X_n\}_{n=1}^{\infty}$ adapted to a filtration $\{\mathcal{F}_n\}_{n=0}^{\infty}$,*

$$\mathbb{E}(X_n|\mathcal{F}_m) = X_m$$

almost surely for all $n \geq m$.

Submartingales and supermartingales are used in the proof of the optimal stopping theorem, so we define them here.

Definition 3.3. Consider a filtration $\{F_n\}_{n=0}^{\infty}$ and a sequence of random variables $\{X_n\}_{n=0}^{\infty}$ adapted to this filtration. $\{X_n\}_{n=0}^{\infty}$ is a submartingale if for each n

$$X_n \leq \mathbb{E}(X_{n+1})|\mathcal{F}_n)$$

and it is a supermartingale if for each n

$$X_n \geq \mathbb{E}(X_{n+1})|\mathcal{F}_n)$$

4. Stopping times

Definition 4.1. A random variable T is a stopping time for a filtration if it takes values in $\{0, 1, 2, \dots\} \cup \{\infty\}$ and for each $n < \infty$, the set $\{T = n\}$ is $\mathcal{F}_n$ measurable.

Intuitively, this means that stopping times can tell us when to stop based only on information that we possess before we have to stop.

From any stopping time T we get a σ-algebra $\mathcal{F}_T$ and a random variable X_T that is $\mathcal{F}_T$ measurable defined as follows:

Definition 4.2. Let T be a stopping time. Then the σ-algebra $\mathcal{F}_T$ is defined as

$$\mathcal{F}_T = \{A \in \mathcal{F} | A \cap \{T = n\} \in \mathcal{F}_n \ \forall \ n\}$$

The intuition behind this definition is that events are in F_T if whether or not they occur can be determined using only information up to time T.

Definition 4.3. For any $\omega \in \Omega$, the stopped random variable X_T is defined as

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

$\mathcal{F}_T$ and X_T both play important roles in the proofs of the optimal and optional stopping theorems.

Proposition 4.4. *For any stopping time T, X_T is $\mathcal{F}_T$ measurable.*

Proof. Let $B \in \mathbb{R}$ be measurable. Then

$$X_T^{-1}(B) \cap \{T = n\} = X_n^{-1}(B) \in \mathcal{F}_n$$

so that $X_T^{-1}(B) \in \mathcal{F}_T$. Thus, X_T is $\mathcal{F}_T$ measurable. □

Proposition 4.5. *Let S and T be stopping times such that $S \leq T$. Then*

$$\mathcal{F}_S \subset \mathcal{F}_T$$

Proof. Let $A \in \mathcal{F}_S$. Then, since $S \leq T$, for all n,

$$A \cap \{T = n\} = A \cap \{S \leq n\} \cap \{T = n\}.$$

Since $A \in \mathcal{F}_S$, $A \cap \{S \leq n\} \in \mathcal{F}_n$. Then, since $\{T = n\} \in \mathcal{F}_n$, $A \cap \{T = n\} \in \mathcal{F}_n$ which completes the proof. □

It is often helpful for stopping times to be bounded (for example, it is a necessary condition of the optional stopping theorem which will be used to prove the main theorem of this paper). The following technique allows us to construct a bounded stopping time out of an unbounded stopping time. For a stopping time T for a filtration $\{F_n\}_{n=0}^{\infty}$, fix any $n \geq 0$. We define $T \wedge n$ by

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

To prove that $T \wedge n$ is a stopping time, consider that for any $k \geq 0$

$$\{T \wedge n > k\} = \begin{cases} \{T > k\} & \text{if } n > k \\ \emptyset & \text{if } n \leq k \end{cases}.$$

For all $j < k$, $\{T = j\} \in \mathcal{F}_j \subset F_k$ so that

$$\{T > k\} = \bigcup_{j=0}^{k-1} \{T = j\} \in \mathcal{F}_k.$$

Furthermore, $\emptyset \in \mathcal{F}_k$. Thus, $\{T \wedge n > k\} \in \mathcal{F}_k$. But, by the same reasoning, $\{T \wedge n > k-1\} \in \mathcal{F}_{k-1} \subset \mathcal{F}_k$. Then

$$\begin{aligned} \{T \wedge n = k\} &= \{T \wedge n > k-1\} \cap \{T \wedge n \leq k\} \\ &= \{T \wedge n > k-1\} \cap \{T \wedge n > k\}^C \in \mathcal{F}_k \end{aligned}$$

so that $T \wedge n$ is a stopping time. With these definitions and basic properties out of the way, we are ready to move on to proving one of our major theorems.

5. Finite horizon optimal stopping theorem

We begin this section with the optional stopping theorem, a related result to the optimal stopping theorem that has applications of its own, but which we will be using in our proof. Intuitively, this theorem says that in a "fair" process (i.e. one with martingales), the expected outcome cannot be changed by stopping at a particular time.

Theorem 5.1. *Let $\{X_n\}_{n=0}^{\infty}$ be a martingale adapted to $\{F_n\}_{n=0}^{\infty}$ and S and T be bounded stopping times for $\{F_n\}_{n=0}^{\infty}$ such that $S \leq T$. Then X_S and X_T are integrable and*

$$\mathbb{E}(X_T|\mathcal{F}_S) = X_S.$$

In particular, we can set $S = 0$ and see that for any bounded stopping time T,

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

Proof. Since S and T are bounded, there exists $n \geq 0$ such that

$$S \leq T \leq n.$$

Then for all $\omega \in \Omega$, there exists $0 \leq i \leq n$ such that $X_T(\omega) = X_i(\omega)$. Then for all $\omega \in \Omega$,

$$|X_T(\omega)| \leq \sum_{i=0}^{\infty} |X_i(\omega)|$$

so that $|X_T|$ is bounded by $\sum_{i=0}^{\infty} |X_i|$ and thus X_T is integrable. A similar argument shows that X_S is integrable.

Consider $A \in \mathcal{F}_S$. Since $T \in \{0, . . ., n\}$, $A = \bigcup_{i=1}^{n} \{T = i\} \cap A$. Then

$$\mathbb{E}(X_n; A) = \sum_{i=0}^{n} \mathbb{E}(X_n; \{T = i\} \cap A).$$

Since $A \in \mathcal{F}_{\mathcal{S}} \subset \mathcal{F}_T$, $A \in \mathcal{F}_T$ so that, by the definition of $\mathcal{F}_T$, $\{T = i\} \cap A \in \mathcal{F}_i$. Furthermore, since $\{X_i\}_{i=0}^{\infty}$ is a martingale, by Proposition 3.2, $\mathbb{E}(X_n|\mathcal{F}_i) = X_i$ for all $i \leq n$. Thus, by the definition of conditional probability,

$$\begin{aligned}\mathbb{E}(X_n; A) &= \sum_{i=0}^{n} \mathbb{E}(X_n; \{T = i\} \cap A) \\ &= \sum_{i=0}^{n} \mathbb{E}(\mathbb{E}(X_n|\mathcal{F}_i); \{T = i\} \cap A) \\ &= \sum_{i=0}^{n} \mathbb{E}(X_i; \{T = i\} \cap A) \\ &= \mathbb{E}(X_T; A).\end{aligned}$$

Replacing all T with S we get that $\mathbb{E}(X_n; A) = \mathbb{E}(X_S; A)$ for all $A \in \mathcal{F}_S$ so that $\mathbb{E}(X_T; A) = \mathbb{E}(X_S; A)$. Therefore, since by Proposition 4.4, X_S is $\mathcal{F}_S$ measurable,

$$\mathbb{E}(X_T|\mathcal{F}_S) = X_S.$$

□

Applying similar reasoning for a supermartingale, we get that $\mathbb{E}(X_T|\mathcal{F}_S) \leq X_S$ almost surely [1].

We now state and prove a lemma about martingales that will prove important in the optimal stopping theorem proof. A similar statement holds for submartingales and supermartingales, but it is not necessary for our proof.

Lemma 5.2. *Let $\{X_n\}_{n=0}^{\infty}$ be a sequence of integrable random variables adapted to the filtration $\{F_n\}_{n=0}^{\infty}$ and let T be a stopping time for this filtration. Assume that for all n, $\mathbb{E}(X_{n+1}|\mathcal{F}_n) = X_n$ almost surely on the set $\{T > n\}$. Then $\{X_{T\wedge n}\}_{n=0}^{\infty}$ is a martingale adapted to $\{F_n\}_{n=0}^{\infty}$.*

Proof. Let $n \geq 0$. Since $T \wedge n \leq n$, by Proposition 4.5 $\mathcal{F}_{T\wedge n} \subset \mathcal{F}_n$. By Proposition 4.4, $X_{T\wedge n}$ is $\mathcal{F}_{T\wedge n}$ measurable, so it is $\mathcal{F}_n$ measurable.

A similar argument to that in the proof of the optional stopping theorem shows that $|X_{T\wedge n}| \leq \sum_{i=0}^{n} |X_i|$ so that $X_{t\wedge n}$ is integrable.

Since $\Omega = \bigcup_{i=0}^{n} \{T = i\} \cup \{T > n\}$,

$$\begin{aligned}\mathbb{E}(X_{T\wedge(n+1)}|\mathcal{F}_n) &= \sum_{i=1}^{n} \mathbb{E}(X_{T\wedge(n+1)}\chi_{\{T=i\}}|\mathcal{F}_n) + \mathbb{E}(X_{T\wedge(n+1)}\chi_{\{T>n\}}|\mathcal{F}_n) \\ &= \sum_{i=1}^{n} \mathbb{E}(X_i\chi_{\{T=i\}}|\mathcal{F}_n) + \mathbb{E}(X_{n+1}\chi_{\{T>n\}}|\mathcal{F}_n)\end{aligned}$$

For $i < n$, $\mathcal{F}_i \subset \mathcal{F}_n$ so X_i is $\mathcal{F}_n$ measurable. Since $\{T = i\} \in \mathcal{F}_i \subset \mathcal{F}_n$, $\chi_{\{T=i\}}$ is $\mathcal{F}_n$ measurable. Furthermore, $\{T > n\} \in \mathcal{F}_n$ so that $\chi_{\{T>n\}}$ is $\mathcal{F}_n$ measurable. Thus, by Propositions 2.7 and 2.8,

$$\mathbb{E}(X_{T\wedge(n+1)}|\mathcal{F}_n) = \sum_{i=0}^{n} X_i\chi_{\{T=i\}} + \chi_{\{T>n\}}\mathbb{E}(X_{n+1}|\mathcal{F}_n)$$

But

$$\sum_{i=0}^{n} X_i \chi_{\{T-i\}} = X_T \chi_{\{T \le n\}}$$

and by assumption,

$$\chi_{\{T>n\}} \mathbb{E}(X_{n+1}|\mathcal{F}_n) = \chi_{\{T>n\}} X_n$$

so that

$$\mathbb{E}(X_{T\wedge(n+1)}|\mathcal{F}_n) = X_T \chi_{\{T\le n\}} + X_n \chi_{\{T>n\}} = X_{T\wedge n}$$

which completes the proof. □

We now move to the optimal stopping theorem. This theorem gives us a stopping time T for a filtration $\{F_n\}_{n=1}^N$ and a sequence of integrable random variables $\{X_n\}_{n=1}^N$ such that $\mathbb{E}(X_T)$ is maximized. In applications, this can be used to execute a decision on when to stop a process that is providing some outcome, signified by the X_n's, in order to maximize the expected value of the outcome when you stop.

We will now derive such a stopping time, and then state and prove the theorem that this stopping time actually maximizes $\mathbb{E}(X_T)$. We will be working over a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with a sequence of integrable random variables $\{X_n\}_{n=1}^\infty$. Define $Y_n := \mathbb{E}(X_n|\mathcal{F}_n)$ for the filtration $\{\mathcal{F}_n\}_{n=1}^\infty$.

Lemma 5.3. *Let T be a stopping time in $\{1, . . ., N\}$ for $\{F_n\}_{n=1}^N$. Then*

$$\mathbb{E}(X_T) = \mathbb{E}(Y_T).$$

Proof. Since $\{T = n\} \in \mathcal{F}_n$, for each n, $\mathbb{E}(Y_n; \{T = n\}) = \mathbb{E}(X_n; \{T = n\})$. Thus,

$$\mathbb{E}(Y_T) = \sum_{n=1}^{N} \mathbb{E}(X_n; \{T = n\}) = \mathbb{E}(X_T)$$

which completes the proof. □

We now define what is called the Snell envelope of $Y_1, . . ., Y_N$ via backwards induction. Define $V_N := Y_N$ and then for each $n < N$,

$$V_n := \max\{Y_n, \mathbb{E}(V_{n+1}|\mathcal{F}_n)\}.$$

Then V_n is a supermartingale adapted to $\{F_n\}_{n=1}^N$ and $V_n \ge Y_n$ almost surely for every n. V_n is actually the smallest supermartingale with this property, but we will not prove or use that statement in this paper.

We now define τ which is the stopping time that will maximize $E(X_\tau)$ as

$$\tau := \min\{n : V_n = Y_n\}.$$

Since $V_N = Y_N$, $\tau \le N$. Since

$$\{\tau = n\} = \{Y_n = V_n\} \cap \{Y_k \ne V_k \forall k < n\}$$

which is the intersection of $\mathcal{F}_n$ measurable sets, $\{\tau = n\}$ is $\mathcal{F}_n$ measurable and τ is a stopping time for $\{\mathcal{F}_n\}_{n=1}^N$. We now state and prove the optimal stopping theorem.

Theorem 5.4. *The finite horizon optimal stopping theorem states that the stopping time τ as defined above maximizes $E(X_\tau)$ with respect to all other stopping times for the filtration $\{\mathcal{F}_n\}_{n=1}^N$, and this maximum value is equal to $\mathbb{E}(V_1)$.*

Proof. Let T be a stopping time for the filtration $\{\mathcal{F}_n\}_{n=1}^N$. When $n < \tau$, $V_n \neq Y_n$ by definition of τ so $V_n = \mathbb{E}(V_{n+1}|\mathcal{F}_n)$. Then by lemma 5.2, $\{V_{\tau\wedge n}\}_{n=1}^N$ is a martingale adapted to $\{\mathcal{F}_n\}_{n=1}^N$. Then

$$\mathbb{E}(Y_\tau) = \mathbb{E}(V_\tau) = \mathbb{E}(V_{\tau\wedge n}) = \mathbb{E}(V_{\tau\wedge 1}) = \mathbb{E}(V_1)$$

where the first equality is by definition of τ, the second equality is because $\tau \leq N$ so that $\tau \wedge N = \tau$, the third equality is by the Theorem 5.1, and the last equality is because $1 \leq \tau$ so that $\tau \wedge 1 = 1$. Then, by the optional stopping theorem for supermartingales and the definition of V_T,

$$\mathbb{E}(X_\tau) = \mathbb{E}(Y_\tau) = \mathbb{E}(V_1) \geq \mathbb{E}(V_T) \geq \mathbb{E}(Y_T) = \mathbb{E}(X_T)$$

□

We now detail two examples of applications of the finite horizon optimal stopping theorem.

Example 5.5. Suppose you are seeking an employee to fill a job opening. You have N candidates to interview, and you must offer the position to a candidate immediately following their interview; if you do not offer them the job immediately, you may not go back and offer it to them later. When should you stop interviewing candidates? This theorem can help make that decision so that you don't hire too early, and miss better candidates later, or hire too late, and leave behind good candidates in the early round.

Start with some $N \geq 2$ number of candidates and denote the rank of the i^{th} candidate among the first n candidates by r_i^n. Denote r_i^N by r_i. We will assume that $r_1, \ldots, r_n$ is uniformly distributed over the set of permutations of $\{1, \ldots, N\}$. Then for each $n \leq N$, $r_1^n, \ldots, r_n^n$ is uniformly distributed over the set of permutations of $1, \ldots n$.

At time n, the only information the interviewer has is $r_1^n, \ldots, r_n^n$. Thus, let $\mathcal{F}_n$ be the σ-algebra generated by the vector-valued random variable $r^n := (r_1^n, \ldots, r_n^n)$. Clearly $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \ldots \subset \mathcal{F}_N$. We want τ to maximize $\mathbb{P}(r_\tau = 1)$ so we want to define $\{X_n\}_{n=1}^n$ such that maximizing $\mathbb{E}(X_\tau)$ also maximizes $\mathbb{P}(X_\tau)$. Thus, define

$$X_n := \chi_{\{r_n=1\}}$$

so that $\mathbb{E}(X_n) = \mathbb{P}(r_n = 1)$.

Step 1: Computing $\{Y_n\}_{n=1}^N$ as defined in the proof of Theorem 5.4, that is

$$Y_n := \mathbb{E}(X_n|\mathcal{F}_n).$$

Since $X_n = \chi_{\{r_n=1\}}$, $\mathbb{E}(X_n|\mathcal{F}_n) = \mathbb{P}(r_n = 1|\mathcal{F}_n)$. If $r_n^n \neq 1$ then $r_n \neq 1$ so we only need to consider $r_n^n = 1$. If this is the case, then $r_1^n, \ldots r_{n-1}^n$ form a permutation of $2, \ldots, n-1$, denoted by $r_1^n = \sigma_1, \ldots, r_{n-1}^{n-1} = \sigma_{n-1}$. We want to compute

$$\mathbb{P}(r_n = 1|r_1^n = \sigma_1, \ldots, r_{n-1}^n = \sigma_{n-1}, r_n^n = 1).$$

If $r_n = 1$ then $r_n^n = 1$ so that

$$\begin{aligned}\mathbb{P}(r_1^n = \sigma_1, \ldots, r_{n-1}^n = \sigma_{n-1}, r_n^n = 1|r_n = 1) &= \mathbb{P}(r_1^n = \sigma_1, \ldots, r_{n-1}^n = \sigma_{n-1}) \\ &= \frac{1}{(n-1)!}.\end{aligned}$$

This allows us to use Bayes' rule to compute the conditional probability we actually want as follows:

$$\begin{aligned}\mathbb{P}(r_n = 1 | r_1^n = \sigma_1, \ldots, r_{n-1}^n = \sigma_{n-1}, r_n^n = 1) &= \frac{\mathbb{P}(r_1^n = \sigma_1, \ldots, r_{n-1}^n = \sigma_{n-1}, r_n^n = 1 | r_n = 1)\mathbb{P}(r_n = 1)}{\mathbb{P}(r_1^n = \sigma_1, \ldots, r_{n-1}^n = \sigma_{n-1}, r_n^n = 1)} \\ &= \frac{\frac{1}{(n-1)!}\frac{1}{N}}{\frac{1}{n!}} \\ &= \frac{n}{N}\end{aligned}.$$

Thus, by Proposition 2.3,

$$Y_n = \mathbb{E}(X_n | \mathcal{F}_n) = \mathbb{P}(r_n = 1 | \mathcal{F}_n) = \frac{n}{N}\chi_{\{r_n^n = 1\}}.$$

Step 2: Calculating V_n's. Let the V_n's be defined as in the proof of the optimal stopping theorem. In order to calculate V_n, it is helpful to show that for each $2 \leq n \leq N$, Y_n is independent of $\mathcal{F}_{n-1}$. To prove this, it suffices to show that $\mathbb{P}(Y_n = \frac{n}{N} | \mathcal{F}_{n-1}) = \mathbb{P}(r_n^n = 1 | \mathcal{F}_n)$ is nonrandom (i.e. a fixed constant). Using a similar strategy as above, we can see that

$$\begin{aligned}\mathbb{P}(r_n^n = 1 | \mathcal{F}_{n-1}) &= \mathbb{P}(r_n^n = 1 | r_1^{n-1} = \sigma_1, \ldots, r_{n-1}^{n-1} = \sigma_{n-1}) \\ &= \frac{\mathbb{P}(r_n^n = 1, r_1^{n-1} = \sigma_1, \ldots, r_{n-1}^{n-1} = \sigma_{n-1})}{\mathbb{P}(r_1^{n-1} = \sigma_1, \ldots, r_{n-1}^{n-1} = \sigma_{n-1})} \\ &= \frac{\mathbb{P}(r_n^n = 1, r_1^n = \sigma_1 + 1, \ldots, r_{n-1}^n = \sigma_{n-1} + 1)}{\mathbb{P}(r_1^{n-1} = \sigma_1, \ldots, r_{n-1}^{n-1} = \sigma_{n-1})} \\ &= \frac{\frac{1}{n!}}{\frac{1}{(n-1)!}} \\ &= \frac{1}{n}\end{aligned}$$

Thus, $\mathbb{E}(Y_n = \frac{n}{N} | \mathcal{F}_n) = \frac{1}{n}$ and Y_n is independent from $\mathcal{F}_{n-1}$.

We now want to use this to show that for each $1 \leq n \leq N-1$, $\mathbb{E}(V_{n+1} | \mathcal{F}_n)$ is equal to a nonrandom quantity, which we will denote by v_n^N. We proceed by backwards induction. Since $V_N = Y_N$, the claim holds for $n = N+1$. Furthermore, $v_{N-1}^N = \frac{1}{N}$. Assume the claim holds for some $1 \leq n \leq N-1$. Then

$$\begin{aligned}\mathbb{E}(V_n | \mathcal{F}_{n-1}) &= \mathbb{E}(\max\{Y_n, \mathbb{E}(V_{n+1} | \mathcal{F}_n)\} | \mathcal{F}_{n-1}) \\ &= \mathbb{E}(\max\{Y_n, v_n^N\} | \mathcal{F}_{n-1}).\end{aligned}$$

Since v_n^N is nonrandom, it is independent of $\mathcal{F}_{n-1}$, and so is Y_n. Thus, $\max\{Y_n, v_n^N\}$ is independent of $\mathcal{F}_{n-1}$ so that $\mathbb{E}(V_n | \mathcal{F}_{n-1}) = \mathbb{E}(\max\{Y_n, v_n^N\} | \mathcal{F}_{n-1})$ is nonrandom. This gives us nonrandom quantities $v_1^N, \ldots, v_{N-1}^N$ such that

$$v_n^N = \mathbb{E}(V_{n+1} | \mathcal{F}_n).$$

As shown above, $v_{N-1}^N = \frac{1}{N}$ and for every $n < N-1$, since $v_{n-1}^N = \mathbb{E}(\max\{Y_n, v_n^N\})$, we have that

$$v_1^N \geq \ldots \geq v_{N-1}^N = \frac{1}{N}.$$

Define

$$t_N := \min\{n \leq N \mid v_n^N \leq \frac{n}{N}\}$$

which exists because $v_{N-1}^N = \frac{1}{N} \leq \frac{N-1}{N}$. Then for $n < t_N$,

$$v_n^N \geq v_{t_N-1}^N > \frac{t_N - 1}{N} > \frac{n}{N}$$

and for $t_N \leq n$,

$$v_n^N \leq v_{t_N}^N \leq \frac{n}{N}.$$

This implies that for $n < t_N$, since $v_n^N > \frac{n}{N} \geq Y_n$,

$$v_{n-1}^N = \mathbb{E}(\max\{Y_n, v_n^N\}|\mathcal{F}_{n-1}) = \mathbb{E}(v_n^N|\mathcal{F}_{n-1}) = \mathbb{E}(v_n^N) = v_n^N.$$

Furthermore, for $t_N \leq n \leq N-1$,

$$\begin{aligned} v_{n-1}^N &= \mathbb{E}(\max\{Y_n, v_n^N\}|\mathcal{F}_{n-1}) \\ &= \frac{n}{N}\mathbb{P}(Y_n = \frac{n}{N}) + v_n^N \mathbb{P}(Y_n = 0) \\ &= \frac{1}{N} + v_n^N(1 - \frac{1}{n}) \end{aligned}$$

We will now use this information to show via backwards induction that for $t_N - 1 \leq n \leq N-1$,

$$v_n^N = \frac{1}{N} + \frac{n}{N}\sum_{k=n+1}^{N-1}\frac{1}{k}.$$

Clearly this holds for $N-1$. Assume this holds for some n. Then

$$\begin{aligned} v_{n-1}^N &= \frac{1}{N} + v_n^N(1 - \frac{1}{n}) \\ &= \frac{1}{N} + \left(\frac{1}{N} + \frac{n}{N}\sum_{k=n+1}^{N-1}\frac{1}{k}\right)(1 - \frac{1}{n}) \\ &= \frac{1}{N} + \frac{1}{N} - \frac{1}{nN} + \frac{n-1}{N}\sum_{k=n+1}^{N-1}\frac{1}{k} \\ &= \frac{1}{N} + \frac{n-1}{N}\left(\frac{1}{n}\right) + \frac{n-1}{N}\sum_{k=n+1}^{N-1}\frac{1}{k} \\ &= \frac{1}{N} + \frac{n-1}{N}\sum_{k=n}^{N-1}\frac{1}{k}. \end{aligned}$$

Thus, t_N can be explicitly defined as the unique integer $n \in \{1, \ldots, N-1\}$ such that

$$\frac{1}{N} + \frac{n-1}{N}\sum_{k=n}^{N}\frac{1}{k} > \frac{n}{N} \geq \frac{1}{N} + \frac{n}{N}\sum_{k=n+1}^{N-1}\frac{1}{k}.$$

For all $n < t_N$,

$$\frac{n}{N} < v_n^N = \mathbb{E}(V_{n+1}|\mathcal{F}_n)$$

but

$$V_n = \max\{Y_n, \mathbb{E}(V_{n+1}|\mathcal{F}_n)\} \geq \mathbb{E}(V_{n+1}|\mathcal{F}_n) > \frac{n}{N} \geq Y_n$$

so that $V_n \neq Y_n$ for all $n < t_N$.

For $n \geq t_N$, if $r_n^n = 1$ then $Y_n = \frac{n}{N}$. Then, because $V_n \leq \frac{n}{N}$,

$$V_n = \max\{Y_n, \mathbb{E}(V_{n+1}|\mathcal{F}_n)\} = \max\{\frac{n}{N}, \mathbb{E}(V_{n+1}|\mathcal{F}_n)\} = \frac{n}{N}.$$

Therefore, the optimal time to hire a candidate is the first candidate after the t_N^{th} candidate (where t_N is just an integer that can be calculated for each N) who is better than all candidates who came before them (i.e. $r_n^n = 1$).

Our next example deals with selling an asset.

Example 5.6. Suppose you possess an asset and are looking for the optimal time to sell it. Let $\{X_n\}_{n=1}^N$ be independently and identically distributed random variables, and let $S_n = \sum_{i=1}^n X_i$ represent the price of an asset on day n. Define $\mathcal{F}_n = \sigma(X_1, \ldots, X_n)$. Let $K > 0$ represent the fee that the you must pay to sell the asset. You may choose to never sell, represented as selling at day ∞. Consider stopping times T in the set $\{1, \ldots, N\} \cup \{\infty\}$ for the filtration $\{\mathcal{F}_n\}_{n=1}^N$. If you stop at time $T < \infty$, you will receive a reward $S_n - K$ and if you stop at time $T = \infty$ you will receive a reward of 0.

In order to bypass the problem of working with an infinite stopping time, we will rework the problem as follows. Define $Y_n := (S_n - K)^+$. Any time that an asset's price S_n is less than the fee K, it would be more advantageous for you to never sell the asset, and incur no profit, than to sell the asset at a loss. Thus, we can interpret all values of $S_n - K$ that are less than 0 as simply 0 and adopt the behavior that any time $S_n - K$ evaluates to something less than 0, we don't sell the asset and incur zero profit. Once we obtain a stopping time T that maximizes $\mathbb{E}((S_n - K)^+)$, the optimal stopping time to maximize $\mathbb{E}(S_n - K)$ is

$$\tau := \begin{cases} T & \text{if } S_T - K > 0 \\ \infty & \text{if } S_T - K \leq 0. \end{cases}$$

We will now define functions $w_n : \mathbb{R} \to \mathbb{R}$ that define the optimal expected reward with N steps left. Define $w_0(x) := (x - K)^+$ and for $n > 0$ define

$$w_n(x) := \max\{(x - K)^+, \mathbb{E}(w_{n-1}(x + X_1))\}.$$

<u>Claim:</u> w_{N-n} is the Snell envelope (i.e. V_n as defined in the proof of Theorem 5.4) of Y_n.

Proof. Clearly

$$V_N = (S_N - K)^+ = w_{N-N}(S_N).$$

Assume that for some n, $V_{n+1} = w_{N-n-1}(S_{+1})$. Then

$$\begin{aligned} V_n &= \max\{Y_n, \mathbb{E}(V_{n+1}|\mathcal{F}_n)\} \\ &= \max\{(S_n - K)^+, \mathbb{E}(w_{N-n-1}(S_{n+1})|\mathcal{F}_n)\} \\ &= \max\{(S_n - K)^+, \mathbb{E}(w_{N-n-1}(S_n + X_{n+1})|\mathcal{F}_n)\}. \end{aligned}$$

Because S_n is $\mathcal{F}_n$ measurable and X_{n+1} is independent of $\mathcal{F}_n$,

$$\mathbb{E}(w_{N-n-1}(S_n + X_{n+1})|\mathcal{F}_n) = \mathbb{E}(w_{N-n-1}(S_n + X_{n+1})).$$

Additionally, because X_{n+1} is identically distributed to X_1,

$$\mathbb{E}(w_{N-n-1}(S_n + X_{n+1})|\mathcal{F}_n) = \mathbb{E}(w_{N-n-1}(S_n + X_1)).$$

Thus,

$$\begin{aligned} V_n &= \max\{(S_n - K)^+, \mathbb{E}(w_{N-n-1}(S_n + X_1)\} \\ &= w_{n-N}(S_n) \end{aligned}$$

□

Thus, w_{n-N} is the Snell envelope of Y_n so that the stopping time that maximizes $\mathbb{E}(S_n - K)^+$ is

$$T := \min\{n \mid (S_n - K)^+ = w_{N-n}(S_n)\}.$$

Then

$$\tau = \begin{cases} T & \text{if } Y_T > 0 \\ \infty & \text{if } Y_T \leq 0 \end{cases}.$$

Since there are only N steps, we can manually calculate w_{N-n} to find the actual time to stop.

6. General optimal stopping theorem

There are many problems that the optimal stopping theorem can help solve where the set of times that the process can be stopped at is not finite. Thus, it is helpful to generalize the optimal stopping theorem beyond the finite horizon case. We will now give an adjusted theorem with a few extra conditions that works in the general case, as well as solve an example applying it. Assume we have reward random variables $\{X_n\}_{n=0}^{\infty}$ and a filtration $\{\mathcal{F}_n\}_{n=0}^{\infty}$.

We will show that the following two assumptions, which trivially hold in the finite horizon example, do not hold in the general infinite case, but if they are assumed to hold, we can construct an optimal stopping theorem around them:

- Assumption 1: $\mathbb{E}(\sup_{n\in\mathbb{N}} X_n) < \infty$
- Assumption 2: $\limsup_{n\to\infty} X_n \leq X_\infty$

We first illustrate some examples demonstrating where one of these assumptions does not hold and there is a problem choosing an optimal stopping time.

Example 6.1. First, an example where Assumption 1 fails.
Let $\{Z_n\}_{n=1}^{\infty}$ be Bernoulli random variables with probability parameter p. Define $X_0 = 0$ and

$$X_n := (2^n - 1)\prod_{i=1}^{n} Z_i.$$

Clearly $\limsup_{n\to\infty} X_n = 0$ so that assumption 2 holds. For each $k = 0, 1, 2, \ldots$, $\sup_{n\in\mathbb{N}} X_n = 2^k - 1$ with a probability of $\frac{1}{2^{k+1}}$ (since each $X_k = 2^k - 1$ if and only if $X_i \neq 0$ for all $i \leq k$ which occurs with a probability of $\frac{1}{2^{k+1}}$). Thus

$$\mathbb{E}(\sup_{n\in\mathbb{N}} X_n) = \sum_{k=0}^{\infty} \frac{2^k - 1}{2^{k+1}} = \infty$$

and assumption 1 is not satisfied.

Assume that for some n, X_n has not had a failure at that time. Then $X_n = 2^n - 1$ but

$$\mathbb{E}(X_{n+1}|\mathcal{F}_n) = \frac{1}{2}(2^{n+1} - 1) > 2^n - 1 = X_n$$

so that it is always optimal to continue. However, if we never stop, our outcome will be 0. Thus, there is no stopping time that optimizes our expected return.

Example 6.2. Now, an example where Assumption 2 fails.
Define $X_0 = 0$, $X_n = 1 - \frac{1}{n}$, and $X_\infty = 0$. Then

$$\mathbb{E}(\sup_{n \in \mathbb{N}} X_n) = 1 < \infty$$

so that Assumption 1 is satisfied, but

$$\limsup_{n \to \infty} X_n = 1 > 0 = X_\infty$$

so that Assumption 2 is not satisfied. And indeed, we cannot find an optimal stopping rule because the longer we wait, the more our expected reward is increased, but if we never stop, our reward is zero.

The above examples clearly show why, without these assumptions, we cannot have an optimal stopping rule, and why the infinite case is more complicated. We now begin to develop the proof of a general optimal stopping theorem that has the additional constraints of Assumptions 1 and 2. But first, we must examine the idea of regular stopping times.

Definition 6.3. A stopping time T for the filtration $\{\mathcal{F}\}_{n=0}^\infty$ is regular if

$$\mathbb{E}(X_T | \mathcal{F}_n) > X_n$$

almost surely on the set $\{T > n\}$.

Intuitively, this means that whenever T says that you should not stop until after n, the conditional expected return for stopping at T based on what has occurred up to n is greater than if we had just stopped at n.

Lemma 6.4. *Assume that Assumption 1 holds. Then for every stopping time for the filtration $\{\mathcal{F}_n\}_{n=0}^\infty$, there exists a stopping time $\bar{T}$ that is regular such that*

$$\mathbb{E}(X_T) = \mathbb{E}(X_{\bar{T}}).$$

Proof. Define $\bar{T} = \min\{n \geq 0 \mid \mathbb{E}(X_T | \mathcal{F}_n) \leq X_n\}$. Clearly $\bar{T}$ is a stopping time such that $\bar{T} \leq T$. On the set $\{\bar{T} = n\}$, we have $\mathbb{E}(X_T | \mathcal{F}_n) \leq X_n$ almost surely and on the set $\{\bar{T} = \infty\}$ we have $X_T = X_{\bar{T}} = X_\infty$ since $\mathbb{E}(X_T | \mathcal{F}_n) > X_n$ for all n. Then

$$\begin{aligned}
\mathbb{E}(Y_{\bar{T}}) &= \sum_{n=0}^{\infty} \mathbb{E}(\chi_{\{\bar{T}=n\}} X_n) \\
&\geq \sum_{n=0}^{\infty} \mathbb{E}(\chi_{\{\bar{T}=n\}} \mathbb{E}(X_T | \mathcal{F}_n)) + \mathbb{E}(\chi_{\{\bar{T}=\infty\}} X_\infty) \\
&= \sum_{n=0}^{\infty} \mathbb{E}(\chi_{\{\bar{T}=n\}} X_T) \\
&= \mathbb{E}(X_T)
\end{aligned}$$

Since conditional expectation is monotone, we have that $\mathbb{E}(X_{\bar{T}} | \mathcal{F}_n) \geq \mathbb{E}(X_T | \mathcal{F}_n)$ almost surely. Then, since $\mathbb{E}(X_T | \mathcal{F}_n) > X_n$ almost surely on $\{\bar{T} > n\}$ by definition of $\bar{T}$, we have

$$\mathbb{E}(X_{\bar{T}} | \mathcal{F}_n) \geq \mathbb{E}(X_T | \mathcal{F}_n) > X_n$$

on $\{\bar{T} > n\}$. Therefore, $\bar{T}$ is regular. □

Lemma 6.5. *Suppose that Assumption 1 holds. If T and $\bar{T}$ are regular stopping times, then $\tau := \max\{T, \bar{T}\}$ is a regular stopping time and*

$$\mathbb{E}(X_\tau) \geq \max\{\mathbb{E}(X_T), \mathbb{E}(X_{\bar{T}})\}.$$

Proof. $\tau = T$ except on sets of the form $\{T - n\} \cap \{\bar{T} > n\}$ for some n in which case

$$\mathbb{E}(X_\tau|\mathcal{F}_n) = \mathbb{E}(X_{\bar{T}}|\mathcal{F}_n)$$

almost surely. Thus,

$$\begin{aligned}
\mathbb{E}(X_\tau) &= \sum_{n=0}^{\infty} \mathbb{E}(\chi_{\{T=n\}} X_\tau) \\
&= \sum_{n=0}^{\infty} \mathbb{E}(\chi_{\{T=n\}} \mathbb{E}(X_\tau|\mathcal{F}_n)) \\
&= \sum_{n=0}^{\infty} \mathbb{E}(\chi_{\{T=n\}\cap\{\bar{T}>n\}} \mathbb{E}(X_\tau|\mathcal{F}_n)) + \sum_{n=0}^{\infty} \mathbb{E}(\chi_{\{T=n\}\cap\{\bar{T}\leq n\}} \mathbb{E}(X_\tau|\mathcal{F}_n)) \\
&= \sum_{n=0}^{\infty} \mathbb{E}(\chi_{\{T=n\}\cap\{\bar{T}>n\}} \mathbb{E}(X_{\bar{T}}|\mathcal{F}_n)) + \sum_{n=0}^{\infty} \mathbb{E}(\chi_{\{T=n\}\cap\{\bar{T}\leq n\}} \mathbb{E}(X_T|\mathcal{F}_n)) \\
&\geq \sum_{n=0}^{\infty} \mathbb{E}(\chi_{\{T=n\}\cap\{\bar{T}>n\}} X_n) + \sum_{n=0}^{\infty} \mathbb{E}(\chi_{\{T=n\}\cap\{\bar{T}\leq n\}} X_n) \\
&= \sum_{n=0}^{\infty} \mathbb{E}(\chi_{\{T=n\}} X_n) = \mathbb{E}(X_T).
\end{aligned}$$

Swapping T and $\bar{T}$ gives us that $\mathbb{E}(X_\tau) \geq \mathbb{E}(X_{\bar{T}})$. Therefore, $\mathbb{E}(X_\tau) \geq \max\{\mathbb{E}(X_T), \mathbb{E}(X_{\bar{T}})\}$

We now show that τ is a regular stopping time. For any n, the set $\{\tau > n\} = \{T > n\} \cup \{\bar{T} > n\}$. On the set $\{T > n\}$ we can repeat the argument above but conditioning the main expectation on $\mathcal{F}_n$ to see that

$$\mathbb{E}(X_\tau|\mathcal{F}_n) \geq \mathbb{E}(X_T|\mathcal{F}_n) > X_n$$

by regularity of T. Similarly, replacing T with $\bar{T}$ in the previous argument, we get that

$$\mathbb{E}(X_\tau|\mathcal{F}_n) \geq \mathbb{E}(X_{\bar{T}}|\mathcal{F}_n) > X_n$$

on $\{\bar{T} > n\}$. Thus, $\mathbb{E}(X_\tau|\mathcal{F}_n) > X_n$ on the set $\{\tau > n\}$. □

We are now ready to prove the general optimal stopping theorem:

Theorem 6.6. *Suppose assumptions 1 and 2 both hold. Then there exists a stopping rule τ such that*

$$\mathbb{E}(X_\tau) = \sup_T \mathbb{E}(X_T).$$

Proof. If $\sup_T \mathbb{E}(X_T) = -\infty$ then $\mathbb{E}(X_T) = -\infty$ for all stopping times T. Thus, we assume that $-\infty < \sup_T \mathbb{E}(X_T) < \infty$. Consider a sequence of stopping times $(\bar{T}_n)$ such that $\mathbb{E}(X_{\bar{T}_n}) \to \sup_T \mathbb{E}(X_T)$. Then let (T_n) be a sequence of stopping times where for each n, T_n is the regularized version of $\bar{T}_n$ as defined in Lemma 6.4. Then $\mathbb{E}(X_{T_n}) = \mathbb{E}(X_{\bar{T}_n})$ so that $\mathbb{E}(X_{T_n}) \to \sup_T \mathbb{E}(X_T)$. Define

$$\tau_i = \max\{T_1, . . ., T_i\}.$$

Then by Lemma 6.5, $\mathbb{E}(X_{\tau_i}) \geq \mathbb{E}(X_{T_i})$ so that $\mathbb{E}(X_{\tau_i}) \to \sup_T \mathbb{E}(X_T)$ as well. Furthermore, τ_i is a monotone increasing sequence of regular stopping times. Thus, it converges to a stopping time $\tau = \sup\{\tau_1, \tau_2, \ldots .\}$. Since (τ_i) is a sequence of integers, it either converges to ∞ or there is some N such that for $i > N$, $\tau_i = k$ for some $k \in \mathbb{N}$. Then, by Assumption 2,

$$\limsup_{i\to\infty} X_{\tau_i} \leq X_\tau$$

almost surely. By Assumption 1, $\mathbb{E}(\sup_n X_n < \infty)$ so that $X_{\tau_i} \leq \sup_n X_n$ is X_{τ_i} is integrable. Then, by Fatou's Lemma,

$$\sup_T \mathbb{E}(X_T) = \limsup_{i\to\infty} \mathbb{E}(X_{\tau_i}) \leq \mathbb{E}(\limsup_{i\to\infty} X_{S_i}) \leq \mathbb{E}(X_\tau) \leq \sup_T \mathbb{E}(X_T)$$

so that $\mathbb{E}(X_\tau) = \sup_T \mathbb{E}(X_T)$. □

This theorem guarantees that an optimal stopping time exists. The principle of optimality tells us more about what the actual stopping time is. We must first, however, define what is called an essential supremum:

Definition 6.7. For a collection of random variables $\{X_\alpha\}_{\alpha\in\mathcal{A}}$, we say that a random variable Z is an essential supremum of this collection and write $Z = \operatorname{ess\,sup}_{\alpha\in\mathcal{A}} X_\alpha$ if

(1) $\mathbb{P}(Z \geq X_\alpha) = 1$ for all $\alpha \in \mathcal{A}$.
(2) If $\bar{Z}$ is any other random variable such that $\mathbb{P}(\bar{Z} \geq X_\alpha) = 1$ for all $\alpha \in \mathcal{A}$, then $\mathbb{P}(\bar{Z} \geq Z) = 1$.

This definition allows us to take the supremum of possibly uncountable sets of random variables without compromising the output being a random variable.

We will not prove the optimality equation, which is what allows us to actually compute the stopping time, but Ferguson [3] gives a proof.

Theorem 6.8. *Let*

$$V_N^* := ess \sup_{N\geq n} \mathbb{E}(Y_N|\mathcal{F}_n)$$

If Assumption 1 holds, then

$$V_n^* = \max\{Y_n, \mathbb{E}(V_{n+1}^*|\mathcal{F}_n)\}$$

and the optimal stopping rule is

$$\tau = \min\{n \geq 0 \mid Y_n = V_n^*\}$$

Note the similarity with the finite horizon problem. Additionally note that Assumption 2 is not required. However, without Assumption 2, the stopping rule given above may not actually be optimal.

We now move on to an example of an application of this general rule. It is very common, in situations where we want to find the most rewarding time to stop a process, that we do not know when the process will end. For example, in the employee example in section 5, it is much more common that you will not know how many total candidates will end up applying to your job. In this example, we will look at offers on a house: if one left one's house on the market indefinitely, it would likely continue to receive offers indefinitely, so we cannot use a finite horizon tool to solve this problem.

Example 6.9. Let $\{X_n\}_{n=0}^{\infty}$ be a sequence of independent and identically distributed random variables with distribution $F(x)$. Each time unit that we keep the house on the market, we incur a cost, c. So the net amount we will receive from an offer accepted at time n is

$$S_n = X_n - nc.$$

Note that this is under the assumption that we cannot recall previous offers. Ferguson [3] gives a proof that this particular problem satisfies Assumptions 1 and 2 if the X_n have a finite second moment, and also details how to solve for the actual stopping rule. This requires many lemmas not relevant to this paper, so we do not reproduce the proof here. However, for uniformly distributed offers, Ferguson showed that the optimal rule would be to accept the first offer higher than $\sup_T \mathbb{E}(S_T)$.

Acknowledgments

It is a pleasure to thank my mentor, Leonardo Bonanno for providing me the direction and resources for this project, as well as providing extensive and thorough help that was critical to my understanding of all of the material I worked with. I thank Peter May and Ewain Gwynne for all of their work in organizing this program and for allowing me to participate in it.

References

[1] Sourav Chatterjee. STAT 310 (MATH 230) Lecture notes. https://souravchatterjee.su.domains/stats310notes.pdf
[2] Rick Durrett. Probability: Theory and Examples. Version 5, 2019. https://sites.math.duke.edu/ rtd/PTE/PTE5_011119.pdf
[3] Thomas S. Ferguson. Optimal Stopping and Applications. UCLA. 2008. http://www.math.ucla.edu/ tom/Stopping/Contents.html