# Stochastic Optimal Stopping: Problem Formulations

Farid AitSahlia

**Article Outline**


**Keywords** Dynamic Programming, Free boundary, Variational Inequalities, Essential Supremum, Martingales

See also:
Stochastic Optimal Stopping: Numerical Methods
Dynamic Programming: Continuous-time Optimal Control
Dynamic Programming: Optimal Control Applications
Financial Optimization
Hamilton–Jacobi–Bellman Equation
Operations Research and Financial Markets
Variational Inequalities

## Introduction

A typical stochastic optimal stopping problem consists of the following optimization:

F. AitSahlia
Warrington College of Business, University of Florida, Gainesville, USA

$$\sup \mathbb{E}\left[f(X_\tau,\tau)\right] \quad \text{s.t.} \quad \tau \in \mathcal{T}\,, \tag{1}$$

where $\{X_t\} \equiv X$ is a stochastic process known as the state process, $\mathbb{E}$ its associated expectation operator, $f$ a function measurable with respect to the probability law induced by $X$, and $\mathcal{T}$ a set of stopping times to be defined shortly. In many applications $f(X_\tau,\tau)$ is interpreted as the gain resulting from stopping at time $\tau$ when the state value is $X_\tau$. It should be noted that some problem formulations call for minimizing the expected value, which can be recast in the maximization above.

An example that has been the subject of great interest in mathematical finance/financial engineering is one with $f(x,t) = e^{-rt}\max(K-x,0)$, where $K > 0$ is given, and $X$ is the geometric Brownian process

$$X_t = X_0 \exp\{(r-\sigma^2/2)t + \sigma B_t\}\,,$$

where $r$ and $\sigma$ are given positive constants and $\{B_t\}$ is a standard Brownian motion started at 0. In finance $f(X_t,t)$ represents the discounted payoff that results from the exercise at time $t$ of a put stock option by its holder who is allowed to sell this stock at share price $K$ when it is traded at price $X_t$. The option holder's problem is to find the best time to exercise this option, thus maximizing its payoff, a problem that is mathematically expressed as (1). As will be made precise soon, this optimal exercise time (or more generally stopping time) must be determined only on the basis of past observations. It should be mentioned here that $\{B_t\}$ can also be considered the state process, instead of $X$.

Note that the payoff function $f$ as expressed in (1) does include payoffs that are path-dependent through the usual introduction of additional variables to render a problem Markovian. For example, still in the financial realm, one may consider the payoffs $e^{-rt}(M_t - X_t)$ or $e^{-rt}\max(K-A_t,0)$ that depend on the maximum process $M_t = \max_{s\le t} X_s$ or the average process $A_t = (1/t)\int_0^t X_s\,\mathrm{d}s$.

Stochastic optimal stopping theory, or optimal stopping as it is customarily known, is a specialized type of the (stochastic) dynamic programming approach devised by Bellman [1] in the 1950s. However, actual optimal stopping problems originated in Wald's work on sequential statistical inference (Wald [6]), where the problem is to determine sequentially the sample size that will decide between two statistical hypotheses. Another application involving optimally stopped sequential sampling is in clinical trials, where is studied under the general label of *bandit problems* ( [2].) Ever since these early days, this field has experienced several developments in both theory and applications as described for example in the book of Peskir and Shiryaev [5]. As artificial intelligence has gained significant traction over the past few years, optimal stopping has also found applications in machine learning to help with the selection of hyperparameters ( [4].)

Optimal stopping problems are generally approached from a probabilistic perspective through martingales and Markov processes. When the underlying process $X$ in (1) is a diffusion, they also lead to free-boundary problems for partial differential equations. Optimal stopping problems are rarely solved in closed-form and numerical methods abound, a topic addressed in a companion entry in this *Encyclopedia*.

## Definitions

This section sets up basic definitions that lead to the notion of stopping time. As mentioned before, the decision to stop at time $t$ must be based only on information available up to $t$. In this respect the concept of information set in the form of filtration is first formally presented, followed by that of stopping time.

- Discrete-time filtration
  Given a probability space $(\Omega, \mathcal{F}, P)$, a discrete-time filtration is a collection $(\mathcal{F}_n)_{n\geq 0}$ where each $\mathcal{F}_n$ is a $\sigma$-algebra of subsets of $\Omega$ such that $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \cdots \mathcal{F}$. $\mathcal{F}_n$ represents the information available up to time $n$. It generally consists of at least the set of events that have been determined by the realized values of $X_t$ up to time $n$. The latter is called the natural filtration of $X$ and is often augmented to form $(\mathcal{F}_n)_{n\geq 0}$.

- Continuous-time filtration
  Here the definition is essentially identical to the previous modulo a technical condition. Given a probability space $(\Omega, \mathcal{F}, P)$, a continuous-time filtration is a collection $(\mathcal{F}_t)_{t\geq 0}$ where each $\mathcal{F}_t$ is a $\sigma$-algebra of subsets of $\Omega$ such that $\mathcal{F}_s \subset \mathcal{F}_t \subset \mathcal{F}$ for $s \leq$ t. As in the discrete-time case, $\mathcal{F}_t$ also represents information up to time $t$. Additionally, it is assumed that each $\mathcal{F}_t$ contains all $P$-null sets in $\mathcal{F}$ and that $(\mathcal{F}_t)_{t\geq 0}$ is right-continuous; i. e., $\mathcal{F}_t = \bigcap_{t\leq s} \mathcal{F}_s$ for all $t \geq 0$.

- Stopping time
  Let $I = \{0, 1, 2, \ldots\}$ and $I = [0, \infty]$ when $X$ is, respectively, a discrete-time process and a continuous-time process. A random variable $\tau \colon \Omega \to I$ is a stopping time if $P\{\tau < \infty\} = 1$ and $\{\tau \leq t\} \in \mathcal{F}_t$ for all $t \geq 0$. Often the set $I$ is bounded and therefore the former condition is obviously true. The latter condition expresses the fact that the decision to stop at time $t$

must be based solely on information up to time $t$. In this case $\tau$ is adapted to the filtration $(\mathcal{F}_t)_{t\geq 0}$.

## Solution Methods

There are generally two approaches to solving (1): one based on probabilistic tools and another on partial differential equations (PDE) techniques. However, both start by using the dynamic programming principle of optimality to derive the so-called Bellman equation. When the interval $I$ is of the form $[0,T]$ or $\{1,2,\ldots,N\}$ define $\mathcal{T}_t$ to be, respectively, the set of stopping times in $[t,T]$ and $\{t,t+1,\ldots,N\}$. When $I$ is infinite, $\mathcal{T}_t$ is defined as the set of stopping times in $I$ that are $\geq t$. Then solving (1) is tantamount to determining

- the value function $V(x,t) = \sup_{\tau\in\mathcal{T}_t} \mathbb{E}[f(X_\tau,\tau)|X_0 = x]$, and
- the optimal stopping time $\tau_t^* = \mathrm{argmax}_{\tau\in\mathcal{T}_t}\mathbb{E}f(X_\tau,\tau)$.

A sufficient condition guaranteeing the finiteness of the expectation in (1) is

$$\mathbb{E}\left[\sup_{t\in I}|f(X_t,t)|\right] < \infty\,,$$

which can in fact be relaxed a number of ways.

(I) **The Probabilistic Approach: Martingales**
When $I = \{0,1,\ldots,N\}$, then by the optimality principle of dynamic programming we can write the recursion

$$V(x,N) = f(x,N) \tag{2}$$

$$V(x,n) = \max\left\{\mathbb{E}[V(X_{n+1},n+1)|X_n = x], f(x,n)\right\}, \quad 0\leq n\leq N-1\,. \tag{3}$$

The solution for the system (2)–(3) induces a sequence of random variables $S_n = V(X_n,n)$ that satisfies the following properties:

(i) $S_n = \max\{\mathbb{E}\left[S_{n+1}|\mathcal{F}_n\right], f(X_n,n)\}$;
(ii) $(S_n)_{k\leq n\leq N}$ is the smallest super-martingale that dominates the gain process $G_n = f(X_n,n)_{k\leq n\leq N}$ (i. e.; $S_n\geq G_n$ P-a.s.);
(iii) the stopping time $\tau_n^* = \inf\{n\leq k\leq N : S_n = G_n\}$ is optimal for $0\leq n\leq N$;

(iv) the stopped sequence $(S_{k\wedge\tau_n^*})_{n\leq k\leq N}$ is a martingale.

We recall here that a discrete-time process $(M_n)_n$ is a martingale with respect to a filtration $(\mathcal{F}_n)_n$ (martingale for short) if $E|M_n| < \infty$ for $n \geq 0$ and $E(M_{n+1}|\mathcal{F}_n) = M_n$, P-a.s., for $n \geq 0$. Correspondingly, $(M_n)_n$ is a super-martingale if $E(M_{n+1}|\mathcal{F}_n) \leq M_n$, P-a.s., $n \geq 0$. The process $(S_n)_{k\leq n\leq N}$ is called the Snell envelope and the above characterization is particularly useful to obtain the value function $V$ through linear programming when the state space is finite (see Çinlar p. 212 in [3]).

The generalization of the above result to the case where $I$ is countably infinite requires that the sequence $S_n = V(X_n, n)$ be characterized differently through the concept of essential supremum below, which generalizes in some sense that of deterministic supremum.

**Essential Supremum.** Let $I$ be an arbitrary set and $(Z_n)_{n\in I}$ be a collection of random variables defined on the same probability space. Then there exists a countable subset $J \subset I$ such that $Z^* = \sup_{n\in J} Z_n$ satisfies

(a) $Z_n \leq Z^*$ P-a.s. for each $n \in I$;
(b) for any other random variable $\tilde{Z}$ such that $Z_n \leq \tilde{Z}$ P-a.s. for each $n \in I$, we have $Z^* \leq \tilde{Z}$ P-a.s.

The random variable $Z^*$ is labeled *essential supremum* and is denoted by $\text{esssup}_{n\in I} Z_n$.
As a consequence, we can now rewrite the above Snell envelope when $I = \{1, 2, \dots, N\}$ as

$$S_n = \text{esssup}_{\tau\in\mathcal{T}_n} \mathrm{E}[f(X_\tau, \tau)|\mathcal{F}_n]\ , \quad n \in I\ , \tag{4}$$

where $\mathcal{T}_n$ is the set of stopping times in $\{n, n+1, \dots, N\}$. When $I$ is countable infinite then $S_n$ is correspondingly defined with $\mathcal{T}_n$ as the set of stopping times in $\{n, n+1, \dots\}$. Similarly, $S_n$ satisfies both conditions (a) and (b) and the optimality property (i) above for all $n \geq 0$.
For the continuous-time case, where $I$ is an interval, the value function for problem (1) is the Snell envelope of the gain process $(f(X_t, t))_t$ defined as

$$S_t = \text{esssup}_{\tau\in\mathcal{T}_t} \mathbb{E}[f(X_\tau, \tau)|\mathcal{F}_t]\ , \quad t \in I\ , \tag{5}$$

where $\mathcal{T}_t$ is the set of stopping times in $[t, T]$ for a finite horizon problem or $[t, \infty)$ otherwise. The Bellman equation in its discrete form (3) is now replaced by

$$V(x,t) \geq \max\{\mathbb{E}[V(X_s,s)|X_t = x], f(X_t,t)\}\ , \\ \text{for } s \geq t.$$

Formulation (1) has cast the problem of optimal stopping in a Markovian framework. This is in fact the most common situation in practice and the set-up is not too restrictive as it mirrors well the generic martingale situation fully described in Peskir and Shiryaev [5].

(II) **The Probabilistic Approach: Markov Property and Stopping Boundary**
When $X$ is a Markov process (in discrete- or continuous-time) with state space $E$ the optimal stopping time is defined as

$$\tau^* = \inf\{t \in I : X_t \in \mathcal{S}\}\ ,$$

where $\mathcal{S}$ is a closed subset in $I \times E$. $\mathcal{S}$ and its complement $\mathcal{C}$ in $I \times E$ are such that

$$V(x,t) > f(x,t) \text{ on } \mathcal{C}\ , \\ V(x,t) = f(x,t) \text{ on } \mathcal{S}\ .$$

$\mathcal{C}$ and $\mathcal{S}$ are respectively called the continuation and stopping regions. The intersection $B$ of their closures is called the stopping boundary. It is time-dependent when $I$ is bounded and time-homogeneous when $I$ is unbounded. When $I$ is countably finite and $E$ is discrete, then $B$ can be obtained through the backward recursion (2)–(3).

(III) **The PDE Approach**
When the state process $X$ is a diffusion the boundary $B$ and the value function $V$ can be obtained by solving a free-boundary problem. Alternatively, when only the value function is of interest then it can be obtained as the solution of a variational inequality. If we let $\mathcal{L}$ the infinitesimal operator associated with $X$, then assuming regularity and differentiability where necessary, the free-boundary problem when $I = [0,\infty)$ is stated as

$$\begin{aligned} &\mathcal{L}V = 0 \quad \text{in} \ \ \mathcal{C}\ , \\ &\frac{\partial V}{\partial x} = \frac{\partial f}{\partial x} \ \text{ on } \ B\ . \end{aligned} \tag{6}$$

The latter condition is called smooth-fit. It is in a sense the condition that characterizes the optimality of a solution $V$ of the PDE (6). When $I = [0,T]$ the free-boundary problem becomes:

$$\begin{aligned} \mathcal{L}V + \frac{\partial V}{\partial t} &= 0 \quad \text{in} \;\; \mathcal{C}\,, \\ V &= f \quad \text{on} \;\; E \times \{T\}\,, \\ \frac{\partial V}{\partial x} &= \frac{\partial f}{\partial x} \;\; \text{on} \;\; B\,. \end{aligned} \tag{7}$$

One way to avoid reference to the free boundary $B$ is through the use of variational equalities. For example, the problem above with finite horizon $T$ can be re-expressed as

$$\begin{aligned} \min\{\mathcal{L}V + \frac{\partial V}{\partial t}, V - f\} = 0\,, &\quad \text{on} \;\; E \times [0,T) \\ V = f\,, &\quad \text{on} \;\; E \times \{T\}. \end{aligned}$$

## Solution Methods

Optimal stopping problem formulations in stochastic environments abound and can be applied in a variety of diverse fields. Their solution methods draw from a diverse set of techniques, including dynamic programming, stochastic calculus, and martingale theory, among others.

## References


1 Bellman R (1957) Dynamic Programming. Princeton University Press, Princeton
2 Berry, D. A., Fristedt, B. (1985) Bandit Problems: Sequential Allocation of Experiments. Chapman and Hall, London, New York.
3 Çinlar E (1975) Introduction to Stochastic Processes. Prentice-Hall, Englewood Cliffs
4 Dai, Z., Yu, H., Low, B. K. H., & Jaillet, P. (2019, May). Bayesian optimization meets Bayesian optimal stopping. Proceedings of the 36th International Conference on Machine Learning,Long Beach,California, (pp. 1496-1506).
5 Peskir G, Shiryaev A (2006) Optimal Stopping and Free-Boundary Problems. Birkhäuser Verlag, Basel
6 Wald A (1950) Statistical Decisions Functions. Wiley, New York; Chapman and Hall, London