Stochastic Optimal Stopping: Numerical Methods

Farid AitSahlia

Article Outline

Keywords Dynamic Programming, Markov Chain, Monte Carlo Simulation, Free Boundary, Finite Difference, Neural Network, Linear Complementarity, Spline Approximation, Integral Equation

See also:
Stochastic Optimal Stopping: Problem Formulations
Dynamic Programming: Continuous-time Optimal Control
Dynamic Programming: Optimal Control Applications
Financial Optimization
Hamilton–Jacobi–Bellman Equation
Approximate Dynamic Programming: Linear Programming-Based Approaches
Operations Research and Financial Markets
Variational Inequalities
Approximate Dynamic Programming: Linear Programming-Based Approaches
Linear Complementarity Problem
Linear Programming

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

Introduction

A significant number of optimal stopping problems of practical interest may only be solved through numerical schemes. As many of them have surfaced in the area of mathematical finance, illustrations drawn from that field will be used to describe some of these numerical approaches. Specifically, consider problems in the expected-value maximization form

$$\sup_{\tau \in \mathcal{T}} \mathbb{E}[f(X_\tau, \tau)] \,, \tag{1}$$

where $\mathcal{T}$ is a set of stopping times, f a measurable function and $\{X_t\}_{t \in I} \equiv X$ a Markov process, where I is a time index set that can be either discrete or continuous (see AitSahlia [1] for additional details).

Under technical conditions for its existence, a solution for (1) consists of

- the value function $V(x,t) = \sup_{\tau \in \mathcal{T}_t} \mathbb{E}[f(X_\tau, \tau)|X_t = x]$, where $\mathcal{T}_t$ is the set of stopping times subsequent to t in $\mathcal{T}$,
- the optimal stopping time $\tau_t^* = \text{argmax}_{\tau \in \mathcal{T}_t} \mathbb{E}[f(X_\tau, \tau)]$.

In this context, with E denoting the state space of X, the set $E \times I$ is partitioned into a closed set $\mathcal{S}$ and its complement $\mathcal{C}$ labeled, respectively, stopping and continuation regions. Then

$$\tau_t^* = \inf\{s \geq t : X_s \in \mathcal{S}\} \,. \tag{2}$$

Discrete-Time Models

When $\mathcal{T} = \{0, 1, \ldots, N\}$ for some given $N < \infty$, the most straightforward numerical device is the backwards recursive dynamic algorithm

$$V(x, N) = f(x, N) \tag{3}$$

$$V(x, n) = \max\{\mathbb{E}[V(X_{n+1}, n+1)|X_n = x], f(x, n)\}, \quad 0 \leq n \leq N-1 \tag{4}$$

- **Monte-Carlo simulation approach:** An issue with the above might be implementing the proper numerical scheme to estimate $E[V(X_{n+1}, n+1)|X_n = x]$, especially in light of the so-called *curse of dimensionality* that makes this algorithm inefficient in high dimensions. There are potentially two remedies to this problem: one, for finite $\mathcal{T}$, based on Monte-Carlo

simulation, and another, for infinite $\mathcal{T}$, based on large-scale linear programming (LP). For the former, an efficient and popular algorithm is that of Longstaff and Schwartz [11], which is now viewed within the wider context of approximate dynamic programming (see also [5, 13]). The basic idea of this algorithm is to use Monte Carlo simulation and least-squares regression to estimate $\mathbb{E}[V(X_{n+1}, n+1)|X_n = x]$.

- **Linear programming approach:** For infinite $\mathcal{T}$, the value function is time-homogeneous when X and f are. In this case the value function solves

$$V(x) = \sup_{\tau \in \mathcal{T}_t} E[f(X_\tau)|X_t = x] \tag{5}$$

for all $t \in \{0, 1, \dots\}$ and may be obtained through a LP algorithm thanks to its Snell envelope characterization ([5]). Assuming a transition matrix P for X and a finite state space, which might be genuine or the result of a truncation, the resulting LP is

$$\text{Minimize} \quad \sum_x V(x)$$

subject to

$$\begin{aligned} V(x) &\geq \sum_y P(x,y)V(y)\,, \\ V(x) &\geq f(x)\,, \\ V(x) &\geq 0\,. \end{aligned}$$

See Çinlar [5] and Dynkin and Yushkevich [6] for proofs and further details.

- **Neural networks approach:** Becker et al. [5] approximate the optimal stopping time recursively through a sequence of 0-1 stopping decisions based on multilayer feedforward neural networks. Their approach is quick to identify approximate stopping times τ^* and τ^{**} such that

$$\mathbb{E}[f(X^*_\tau, \tau^*)] \leq \sup_{\tau \in \mathcal{T}} \mathbb{E}[f(X_\tau, \tau)] \leq \mathbb{E}[f(X^{**}_\tau, \tau^{**})]$$

Continuous-Time Models

When both X and its time index I are continuous, there are a number of numerical schemes to generate solutions for (1). Overall, they approximate

either the underlying diffusion process X by a discrete version or the value function and its derivatives in its characterizing expression (e. g., integral representation, partial differential equation.)

- **Weak-convergence approximation approach:** The most general scheme concerning this approach is to approximate the infinitesimal operator $\mathcal{L}$ of X in the free-boundary problem that characterizes the solution of (1). For example, a finite-differences approximation of derivatives in the free-boundary problem

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

  leads to the formulation of an optimal stopping problem for a Markov chain (see Kushner and Dupuis [8]).

  If the process X is explicitly expressed in terms of Brownian motion, then random-walk approximations can directly be used on the latter. This is a fairly well understood procedure for which rates of convergence have been developed (see Lamberton [10]).

- **Integral equation approach:** In this scheme, one makes use of the Doob–Meyer decomposition formula for submartingales (see Karatzas and Shreve [8]) to express the value function V in terms of the boundary, which itself solves a related integral equation. For example, consider a case in American option pricing , with horizon T, payoff function $f(x,t) = e^{-rt}\max(K-x,0)$, and $X_t = X_0 \exp\{(r-\sigma^2/2)t + \sigma W_t\}$, where $K > 0$, $r > 0$, and $\sigma > 0$ are given and $\{W_t\}_t$ is a standard Brownian motion started at 0. Then the value function V can be decomposed as

$$V(x,t) = U(x,t) + \int_t^T [rK\Phi(-d(X, B(t), \tau - t))]\mathrm{d}\tau\,, \tag{6}$$

  where Φ is the cumulative standard normal distribution function, $d(x,y,\tau) = (\ln(x/y) + (r+\sigma^2/2)\tau)/\sigma\sqrt{\tau} - \sigma\sqrt{\tau}$, and $U(x,t) = Ke^{(T-t)}\Phi(-d(x,K,T-t))$. This formula requires the knowledge of the boundary $B = \partial\mathcal{S}$, where $\mathcal{S}$ is the stopping region that identifies the optimal stopping time (2), and which is obtained as the solution of the integral equation

$$K - B(t) = U(B(t), t) + \int_t^T rK\Phi\left(-d(B(t), B(t), \tau - t\right) \mathrm{d}\tau \,. \qquad (7)$$

Efficient and accurate spline approximations of B can be found in Ait-Sahlia and Lai [2].

- **Linear complementarity approach:** An alternative that does not require the explicit determination of the optimal stopping boundary relies on the variational inequality formulation is, in the finite time horizon case:

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

Finite-difference approximations then lead to a linear complementarity problem (see Huang and Pang [7] and Wilmott et al. [14]).

- **Deep neural networks** Sirigano and Spiliopoulos [12] use deep neural networks to approximate the solution to the free boundary problem. Their approach avoids the use of meshes, which quickly reach the curse of dimensionality, but instead rely on random samples of time and space along with stochastic gradient descent.

Conclusions

Stochastic optimal stopping problems are rarely solved in closed form. Their solutions therefore call for numerical techniques, which include finite-difference approximations to solve PDE–related formulations, spline approximations to the optimal stopping boundary, linear complementarity approaches to solving variational inequalities and methods based on linear programming.

References

1 AitSahlia F Stochastic optimal stopping: problem formulations, this Encyclopedia
2 AitSahlia F, Lai TL (2001) Exercise boundaries and efficient approximations to American option prices and hedge parameters. J Comput Finance 4:85–103
3 Becker, S., Cheridito, P. and Jentzen, A., 2019. Deep optimal stopping. The Journal of Machine Learning Research, 20(1), pp.2712-2736.

4 Carrière J (1996) Valuation of early-exercice price of options using simulations and nonparametric regression. Insurance Math Econ 19:19–30
5 Çinlar E (1975) Introduction to Stochastic Processes. Prentice-Hall, Englewood Cliffs
6 Dynkin EB, Yushkevich AA (1969) Markov Processes: Theorems and Problems. Plenum Press, New York
7 Huang J, Pang J-S (1998) Option Pricing and Linear Complementarity. J Comput Finance 2:31–60
8 Karatzas I, Shreve SE (1991) Brownian Motion and Stochastic Calculus. Springer, New York
9 Kushner HJ, Dupuis P (2001) Numerical Methods for Stochastic Control Problems in Continuous Time, 2nd edn. Springer, New York
10 Lamberton D (2002) Brownian optimal stopping and random walks. Appl Math Optim 45:283–324
11 Longstaff FA, Schwartz ES (2001) Valuing American options by simulation: a simple least-squares approach. Rev Financial Stud 14:113–148
12 Sirignano, J. and Spiliopoulos, K., 2018. DGM: A deep learning algorithm for solving partial differential equations. Journal of computational physics, 375, pp.1339-1364.
13 Tsitsiklis JN, Van Roy B (2001) Regression methods for pricing complex American-style options. IEEE Trans Neural Networks 12(4):694–703
14 Wilmott P, Dewynne JN, Howison S (1993) Option Pricing: Mathematical Models and Computation. Oxford Financial Press, Oxford