Optimal Stopping and the Snell Envelope

zamrik.com Research Notes

Dr. T Zamrik

03 February 2025

1. Abstract

We study the optimal stopping problem in continuous time, where an agent chooses a stopping time to maximise the expected value of a payoff process, following the classical framework of Snell (1952) and its continuous-time extension via the Doob–Meyer decomposition. The value function is characterised as the Snell envelope — the smallest supermartingale dominating the payoff — whose generator satisfies a Hamilton–Jacobi–Bellman variational inequality of obstacle type. The optimal stopping time is the first entry into the stopping region, where the value function equals the payoff, and the continuation region is determined by the strict inequality $V > g$. As a canonical application, we solve the perpetual American put option, obtaining the closed-form exercise boundary and value function, and illustrate how the exercise boundary moves with time-to-expiry under the finite-horizon formulation.

2. Introduction

Optimal stopping is one of the oldest problems in stochastic control. The agent observes a stochastic process and must decide when to stop to collect a reward, trading off the option value of waiting against the risk of deterioration. Applications span finance (American options, optimal liquidation), statistics (sequential testing, the secretary problem), economics (investment under uncertainty), and operations research (optimal search).

The continuous-time formulation, due to Snell (1952) in discrete time and extended by El Karoui (1981) and others, connects the stopping problem to the theory of supermartingales. The key object is the **Snell envelope** V_t, which equals the supremum of expected future payoffs over all stopping times after t. It is simultaneously the value function of the stopping problem and the smallest supermartingale that dominates the payoff process $g(X_t)$.

In a Markovian setting, the Snell envelope reduces to a deterministic function $V(x,t)$ of the state and time, characterised as the unique solution to an obstacle problem — a variational inequality that combines the Black–Scholes-type PDE in the continuation region with the constraint $V \geq g$ everywhere. The boundary between the continuation and stopping regions is a free boundary whose location is determined as part of the solution.

This note develops the theory systematically: Section 2 states the problem, Section 3 characterises the Snell envelope, Section 4 derives the HJB variational inequality, Section 5 solves the perpetual American put in closed form, Section 6 presents the finite-horizon formulation and numerical approach, Section 7 displays numerical results, and Section 8 concludes.

3. Problem Formulation

3.1 Stochastic Basis

Fix a filtered probability space $(\Omega, \mathscr{F}, (\mathscr{F}_t)_{t\geq 0}, \mathbb{P})$ satisfying the usual conditions. Let W_t be a standard $\mathscr{F}_t$-Brownian motion and let the state process X_t follow the Itô SDE:

$$dX_t = \mu(X_t, t)\, dt + \sigma(X_t, t)\, dW_t, \qquad X_0 = x. \tag{3.1}$$

3.2 Payoff and Stopping Times

Let $g : \mathbb{R} \times [0, T] \to \mathbb{R}_+$ be a continuous payoff function. Denote by $\mathscr{T}_{t,T}$ the set of $\mathscr{F}$-stopping times taking values in $[t, T]$. The agent's **value function** is:

$$V(x, t) = \sup_{\tau \in \mathscr{T}_{t,T}} \mathbb{E}\left[e^{-r(\tau - t)} g(X_\tau, \tau) \,\middle|\, X_t = x\right], \tag{3.2}$$

where $r \geq 0$ is the discount rate.

The **stopping region** and **continuation region** are:

$$\mathscr{S} = \{(x, t) : V(x, t) = g(x, t)\}, \qquad \mathscr{C} = \{(x, t) : V(x, t) > g(x, t)\}. \tag{3.3}$$

3.3 The Snell Envelope

In the Markovian setting, define the process $\hat{V}_t = V(X_t, t)$. The **Snell envelope** characterisation states:

> *$V(x, t)$ is the smallest function such that (i) $V \geq g$ everywhere, and (ii) $e^{-rt} V(X_t, t)$ is a supermartingale.*

The Doob–Meyer decomposition of the discounted Snell envelope gives:

$$e^{-rt} V(X_t, t) = \mathscr{M}_t - \mathscr{A}_t, \tag{3.4}$$

where $\mathscr{M}_t$ is a martingale and $\mathscr{A}_t$ is an increasing process that grows only on $\mathscr{S}$. This decomposition encodes the fact that it is never optimal to stop in $\mathscr{C}$: the envelope is a local martingale in the continuation region.

4. HJB Variational Inequality

4.1 Derivation

Applying Itô's formula to $e^{-rt}V(X_t, t)$ and using the supermartingale property yields the **HJB variational inequality**:

$$\min\{-\mathscr{L}V + rV,\ V - g\} = 0, \tag{4.1}$$

where $\mathscr{L}$ is the generator of X:

$$\mathscr{L}V = \partial_t V + \mu(x,t)\,\partial_x V + \frac{1}{2}\sigma^2(x,t)\,\partial_{xx}V. \tag{4.2}$$

This is an **obstacle problem**: in the continuation region $\mathscr{C}$, $V > g$ so the minimum is attained by the first term, giving $\mathscr{L}V = rV$ (the PDE is active). In the stopping region $\mathscr{S}$, $V = g$ so the constraint is binding.

4.2 Boundary Conditions

For a finite horizon T: - **Terminal**: $V(x,T) = g(x,T)$ for all x. - **Far field**: $V(x,t) \to 0$ as $x \to 0$ (for put-type payoffs) or $V(x,t) \sim x$ as $x \to \infty$. - **Free boundary**: at the exercise boundary $b(t)$, the **smooth pasting** conditions hold:

$$V(b(t),t) = g(b(t),t), \qquad \partial_x V(b(t),t) = \partial_x g(b(t),t). \tag{4.3}$$

Smooth pasting is the first-order optimality condition ensuring the agent is indifferent between stopping and continuing at the boundary.

5. Perpetual American Put

5.1 Setup

Set $g(x) = (K-x)^+$, $r > 0$, GBM dynamics $dX_t = rX_t\,dt + \sigma X_t\,dW_t$, and $T = \infty$. The time-homogeneity collapses the problem to a one-dimensional ODE. The value function depends only on x, satisfying:

$$\frac{1}{2}\sigma^2 x^2 V''(x) + rx\,V'(x) - rV(x) = 0 \quad \text{for } x > b^*, \tag{5.1}$$

with boundary conditions $V(b^*) = K - b^*$ and $V'(b^*) = -1$ (smooth pasting), and $V(x) \to 0$ as $x \to \infty$.

5.2 Closed-Form Solution

The general solution to the ODE is $V(x) = Ax^\beta + Bx^\gamma$ where β, γ are the roots of the characteristic equation $\frac{1}{2}\sigma^2\alpha(\alpha-1) + r\alpha - r = 0$. The relevant root is:

$$\beta = -\frac{2r}{\sigma^2} < 0, \tag{5.2}$$

and regularity at infinity forces $B = 0$. Applying the smooth-pasting conditions gives:

$$b^* = \frac{\beta}{\beta - 1}K, \qquad V(x) = (K - b^*)\left(\frac{x}{b^*}\right)^{\beta} \quad \text{for } x \geq b^*. \tag{5.3}$$

For $x < b^*$: $V(x) = K - x$ (immediate exercise is optimal).

The exercise boundary $b^* < K$ is always strictly below the strike, reflecting the option value of waiting. As $r \to \infty$, $\beta \to -\infty$ and $b^* \to K$ (discounting dominates). As $\sigma \to \infty$, $\beta \to 0$ and $b^* \to 0$ (uncertainty makes waiting worthless).

6. Finite-Horizon Formulation

6.1 Black–Scholes Obstacle Problem

For the finite-horizon American put with maturity T, the value function $V(x, \tau)$ (where $\tau = T - t$ is time to expiry) solves:

$$\min\left\{-\partial_\tau V + \frac{1}{2}\sigma^2 x^2 \partial_{xx} V + rx\,\partial_x V - rV,\ V - (K - x)^+\right\} = 0, \tag{6.1}$$

with $V(x, 0) = (K - x)^+$.

There is no closed-form solution. The exercise boundary $b(\tau)$ is a decreasing function of τ (i.e., increasing in time-to-expiry), starting at $b(0) = K$ and converging to b^* as $\tau \to \infty$.

6.2 Integral Equation for the Boundary

The exercise boundary satisfies the early-exercise premium decomposition:

$$V(x, \tau) = V^{\mathrm{Eur}}(x, \tau) + \int_0^\tau rK\, e^{-rs}\, N(-d_2(x, b(\tau - s), s))\, ds, \tag{6.2}$$

where V^{Eur} is the European put price and $N(\cdot)$ is the standard normal CDF. This implicit equation for $b(\tau)$ must be solved numerically.

6.3 Numerical Algorithm

```
Input:  sigma, r, K, T, Nx, Nt
Output: V(x, tau), b(tau)

Grid: x_i = linspace(x_lo, x_hi, Nx),  tau_j = linspace(0, T, Nt)
IC:   V[:, 0] = max(K - x, 0)

For j = 1 to Nt:
    dtau = T / Nt
    Solve tridiagonal system (Crank-Nicolson):
```

```
        [I - 0.5*dtau*L] V_new = [I + 0.5*dtau*L] V_old
    where L is the BS operator on x-grid (central differences)
    Apply obstacle:
        V_new[i] = max(V_new[i], K - x[i])   for all i
    Locate boundary:
        b[j] = max x_i such that V[i, j] == K - x[i]

Return V, b
```

7. Numerical Results

7.1 Snell Envelope vs European Price

The Snell envelope $V(x, \tau)$ dominates the European put price everywhere. The gap (early-exercise premium) is concentrated near the exercise boundary and grows with τ.

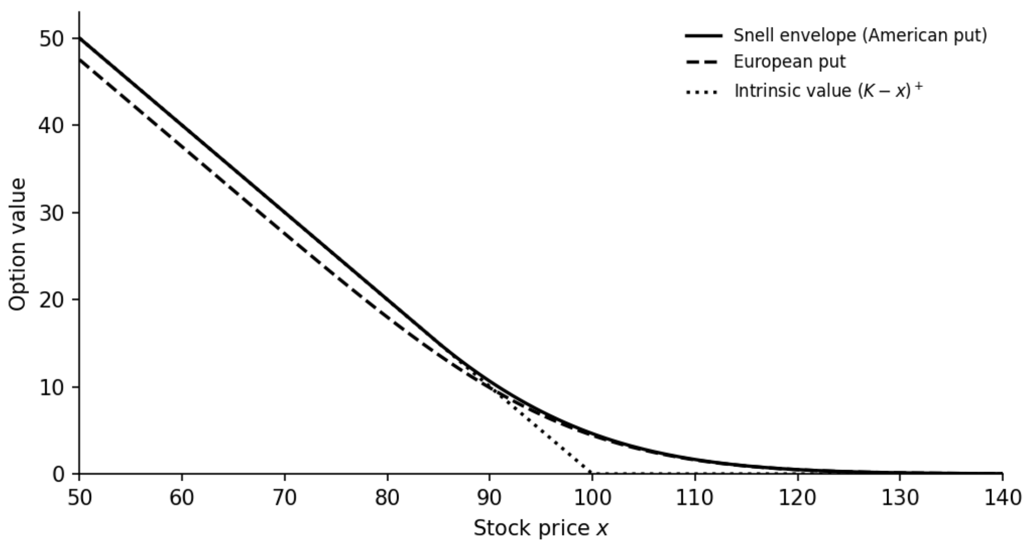

Figure 1: Snell envelope (solid), European put (dashed), and intrinsic value (dotted) as functions of stock price x at $\tau = 0.5$. Parameters: $K = 100$, $r = 0.05$, $\sigma = 0.2$. The early-exercise premium is the vertical gap between the Snell envelope and the European price. The exercise boundary b^* is where the Snell envelope meets the intrinsic value.

7.2 Value Function Across Time-to-Expiry

As τ decreases toward zero, the value function converges to the intrinsic value and the continuation region shrinks.

7.3 Exercise Boundary

The exercise boundary $b(\tau)$ is a decreasing function of time-to-expiry, connecting $b(0) = K$ at expiry to the perpetual boundary b^* at infinite horizon.

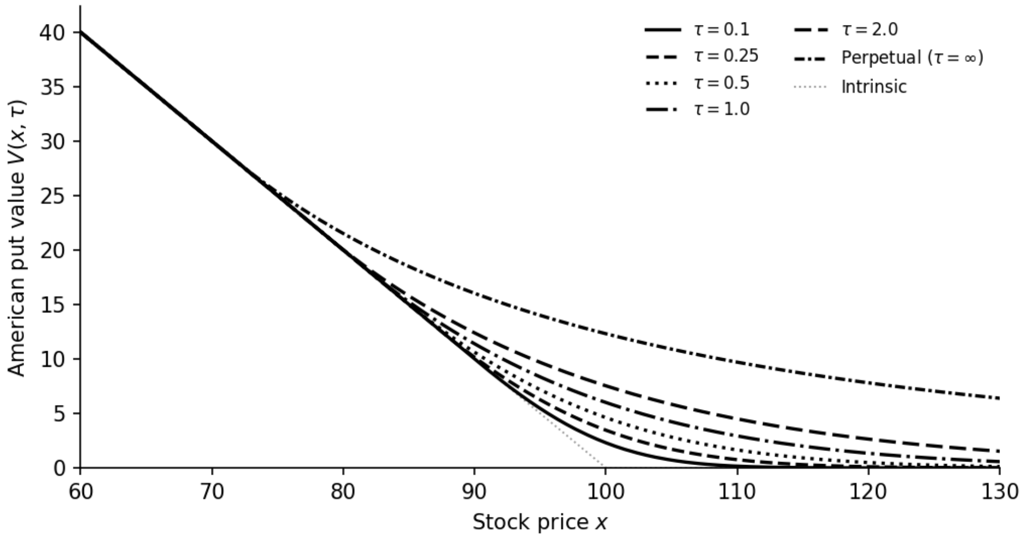

Figure 2: Value function $V(x, \tau)$ for five time-to-expiry levels $\tau \in \{0.1, 0.25, 0.5, 1.0, 2.0\}$, shown as curves of decreasing smoothness. Parameters: $K = 100$, $r = 0.05$, $\sigma = 0.2$. The perpetual American put value is also shown (dash-dotted), approached as $\tau \to \infty$.

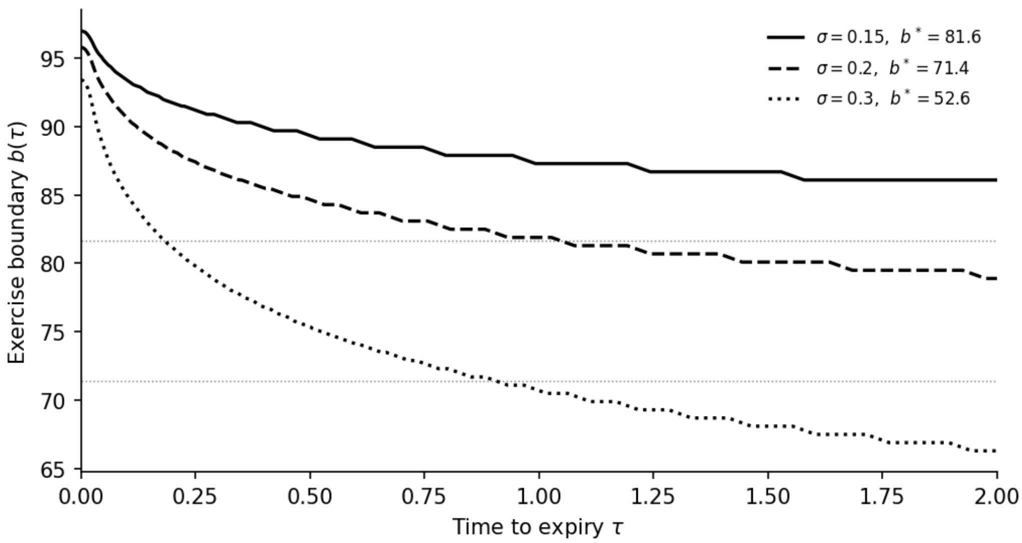

Figure 3: Exercise boundary $b(\tau)$ as a function of time-to-expiry τ for three volatility levels $\sigma \in \{0.15, 0.20, 0.30\}$. Higher volatility lowers the boundary (it is less optimal to exercise early when uncertainty is high). The horizontal asymptote for each curve is the corresponding perpetual boundary $b^* = \frac{\beta}{\beta-1}K$.

8. Conclusion

The Snell envelope provides the canonical unifying framework for optimal stopping: it is simultaneously the value function of the stopping problem, the smallest dominating supermartingale, and the solution to an obstacle PDE. The key structural results are: (i) the value function satisfies a variational inequality that is a PDE in the continuation region and an algebraic constraint in the stopping region; (ii) smooth pasting pins the free boundary; (iii) for the perpetual American put the boundary and value function are available in closed form; (iv) for finite horizon, the boundary must be obtained numerically and satisfies an integral equation involving the early-exercise premium.

Extensions of interest include multi-asset stopping problems (basket options), optimal stopping under model uncertainty (robust Snell envelopes), Lévy process dynamics (replacing the Itô generator with a non-local operator), and sequential problems where the agent can restart after stopping (optimal switching).

9. References

1. Snell, J. L. (1952). Applications of martingale system theorems. *Transactions of the American Mathematical Society*, 73(2), 293–312.
2. El Karoui, N. (1981). Les aspects probabilistes du contrôle stochastique. *Lecture Notes in Mathematics*, 876, 73–238.
3. McKean, H. P. (1965). Appendix: A free boundary problem for the heat equation arising from a problem in mathematical economics. *Industrial Management Review*, 6, 32–39.
4. Samuelson, P. A. (1965). Rational theory of warrant pricing. *Industrial Management Review*, 6, 13–31.
5. Kim, I. J. (1990). The analytic valuation of American options. *Review of Financial Studies*, 3(4), 547–572.
6. Carr, P., Jarrow, R., & Myneni, R. (1992). Alternative characterizations of American put options. *Mathematical Finance*, 2(2), 87–106.
7. Peskir, G., & Shiryaev, A. N. (2006). *Optimal Stopping and Free-Boundary Problems*. Birkhäuser.
8. Øksendal, B. (2003). *Stochastic Differential Equations* (6th ed.). Springer.
9. Lamberton, D., & Lapeyre, B. (2007). *Introduction to Stochastic Calculus Applied to Finance* (2nd ed.). Chapman & Hall.
10. Bensoussan, A., & Lions, J. L. (1982). *Applications of Variational Inequalities in Stochastic Control*. North-Holland.