



Russian Options: HJB and Reflected BSDE — Two Methods, One Price Surface

T. Zamrik

10-FEB-2024

1. Abstract

The Russian option, introduced by Shepp and Shiryaev (1993), is a perpetual American lookback contract whose payoff is the running maximum of the underlying asset. We price it via two independent frameworks: a Hamilton-Jacobi-Bellman (HJB) variational inequality and a singly reflected backward stochastic differential equation (BSDE) with barrier M_t . The HJB approach yields a closed-form solution through a scalar Euler ODE, while the BSDE Bellman iteration — solved by Gauss-Hermite quadrature — converges to the same value function with mean relative error 0.10%. A third result ties the two together: the expected stopping time $u(x) = \mathbb{E}[\tau^* | X_0 = x]$ satisfies a companion Poisson equation $\mathcal{L}u = -1$ under the same operator and boundary conditions, and its closed form reveals that the option premium equals, to leading order, the discount rate times the expected wait times the current value.

2. Introduction

The Russian option is one of the most elegant perpetual optimal stopping problems in mathematical finance. Introduced by Shepp and Shiryaev [1], it pays the holder the running maximum M_τ of the stock price at a stopping time τ , discounted at rate λ . Unlike the American put, which lives in a one-dimensional state space, the Russian option requires tracking both the current price S_t and the running maximum M_t , making the state space two-dimensional.

Two deep theoretical frameworks characterise the value function $V(s, m)$. The first is the HJB theory: V solves a variational inequality in (s, m) space with a Neumann reflection condition at the diagonal $s = m$. By dimensional homogeneity, $V(s, m) = m v(s/m)$, reducing the problem to a scalar Euler ODE in $x = s/m$ with a closed-form solution due to Shepp and Shiryaev. The second is the theory of reflected BSDEs: V is the initial value of a singly reflected BSDE whose Skorokhod regulator is driven by the local time of $M_t - S_t$ at zero — the same local time that generates the Neumann condition in the HJB formulation.

This note develops both frameworks in parallel, demonstrates their numerical agreement, and derives a third result that gives the paper its economic content: the expected stopping time $u(x) = \mathbb{E}[\tau^* | X_0 = x]$ solves the Poisson equation $\mathcal{L}u = -1$ under the same



differential operator \mathcal{L} and the same boundary conditions as the value function. The closed-form solution of this ODE shows that at $x = 1$ (stock at an all-time high), the holder waits approximately 1.84 years before exercising — and that this waiting time is precisely what the premium $V(s, m) - m$ monetises.

3. The Russian Option: Problem Setup

Definition 3.1. (*Filtered probability space.*) Let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{Q})$ be a filtered probability space supporting a standard Brownian motion W_t , where \mathbb{Q} is the risk-neutral measure and $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ is the augmented natural filtration of W .

Under \mathbb{Q} , the asset price follows geometric Brownian motion:

$$dS_t = rS_t dt + \sigma S_t dW_t, \quad S_0 = s > 0, \quad (3.1)$$

where $r > 0$ is the risk-free rate and $\sigma > 0$ is the volatility.

Definition 3.2. (*Running maximum.*) The running maximum process is defined by

$$M_t = \max\left(m, \sup_{0 \leq u \leq t} S_u\right), \quad M_0 = m \geq s.$$

Definition 3.3. (*Russian option value.*) The value of the Russian option with discount rate $\lambda > r$ is

$$V(s, m) = \sup_{\tau \in \mathcal{T}} \mathbb{E}^{\mathbb{Q}}\left[e^{-\lambda\tau} M_\tau \mid S_0 = s, M_0 = m\right],$$

where \mathcal{T} denotes the class of \mathbb{F} -stopping times.

The admissible state space is $\mathcal{D} = \{(s, m) : 0 < s \leq m\}$. The payoff surface is $g(s, m) = m$, independent of s . The running maximum M_t increases only on $\{S_t = M_t\}$, the set of new all-time highs.

Remark 3.4. The condition $\lambda > r$ ensures V is finite for all $(s, m) \in \mathcal{D}$: discounting dominates the drift of M_t .

4. The HJB Variational Inequality

4.1 Derivation

Let \mathcal{L} denote the infinitesimal generator of S_t :

$$\mathcal{L}f(s) = rs f_s(s) + \frac{1}{2}\sigma^2 s^2 f_{ss}(s). \quad (4.1)$$

Theorem 4.1. (*HJB characterisation.*) The value function $V(s, m)$ satisfies the variational inequality

$$\max(m - V, (\mathcal{L} - \lambda)V) = 0 \quad \text{in } \mathcal{D},$$



together with the Neumann reflection condition at the diagonal $s = m$:

$$\frac{\partial V}{\partial m}(m, m) = 0.$$

Proof. Standard dynamic programming on the augmented state (S_t, M_t) ; see Shepp and Shiryaev [1].

The Neumann condition arises because M_t increases only when $S_t = M_t$, so the flux of V in the m -direction at $s = m$ must vanish.

4.2 Reduction to an ODE

By dimensional homogeneity $V(\lambda s, \lambda m) = \lambda V(s, m)$, one writes $V(s, m) = m \cdot v(x)$ with $x = s/m \in (0, 1]$. Substituting into the HJB equation in the continuation region $\{V > m\}$ yields the Euler ODE:

$$\frac{1}{2}\sigma^2 x^2 v''(x) + rx v'(x) - \lambda v(x) = 0, \quad x \in (x^*, 1]. \quad (4.2)$$

The general solution is $v(x) = Ax^{\beta_1} + Bx^{\beta_2}$, where $\beta_1 > 1 > 0 > \beta_2$ solve the characteristic equation:

$$\frac{1}{2}\sigma^2 \beta(\beta - 1) + r\beta - \lambda = 0. \quad (4.3)$$

Lemma 4.2. (Boundary conditions.) The unknowns A , B , and $x^* \in (0, 1)$ are determined by three conditions: 1. **Value match:** $v(x^*) = 1$. 2. **Smooth fit:** $v'(x^*) = 0$. 3. **Neumann at $x = 1$:** $v(1) - v'(1) = 0$. **Proof.** Condition 1 is continuity of V at the stopping boundary. Condition 2 is the smooth-pasting principle [2]. Condition 3 is the image of $\partial_m V(m, m) = 0$ under $x = s/m$.

Proposition 4.3. (Closed-form solution.) The Neumann condition gives $B = -A(1 - \beta_1)/(1 - \beta_2)$. The threshold x^* is the unique solution in $(0, 1)$ of

$$\beta_1 - \frac{(1 - \beta_1)\beta_2}{1 - \beta_2} (x^*)^{\beta_2 - \beta_1} = 0.$$

The constant A follows from value-matching, and $B = -(1 - \beta_1)/(1 - \beta_2) \cdot A$.

The full price surface is:

$$V_{\text{HJB}}(s, m) = \begin{cases} m & \text{if } s/m \leq x^*, \\ m(A(s/m)^{\beta_1} + B(s/m)^{\beta_2}) & \text{if } s/m > x^*. \end{cases} \quad (4.4)$$



5. The Reflected BSDE

5.1 Formulation

Definition 5.1. (*Reflected BSDE.*) A singly reflected BSDE with barrier (L_t) is a triple (Y_t, Z_t, K_t) satisfying

$$Y_t = \xi + \int_t^T f(s, Y_s, Z_s) ds + K_T - K_t - \int_t^T Z_s dW_s,$$

where ξ is the terminal condition, f is the driver, K_t is a non-decreasing Skorokhod regulator, and the flat-off condition holds:

$$Y_t \geq L_t \quad \text{and} \quad \int_0^T (Y_t - L_t) dK_t = 0.$$

For the Russian option: driver $f = -\lambda Y_t$, barrier $L_t = M_t$, terminal condition $\xi = M_T$. Taking $T \rightarrow \infty$ in the discounted formulation, the perpetual reflected BSDE characterises V via its initial value.

Theorem 5.2. (*BSDE characterisation.*) The value function satisfies $V(s, m) = Y_0$, where (Y_t, Z_t, K_t) solves the reflected BSDE with barrier $L_t = M_t$, driver $f = -\lambda Y_t$, and Skorokhod condition $\int_0^\infty (Y_t - M_t) dK_t = 0$. **Proof.** Feynman-Kac theorem for reflected BSDEs [3]; see also El Karoui et al. [4].

5.2 The Skorokhod Condition and Local Time

The flat-off condition forces dK_t to act only on $\{Y_t = M_t\}$. Since M_t increases only when $S_t = M_t$, the regulator is driven by the local time ℓ_t^0 of $M_t - S_t$ at zero:

$$K_t = \int_0^t \mathbf{1}_{\{S_u = M_u\}} d\ell_u^0(M - S). \quad (5.1)$$

Proposition 5.3. (*Local time bridge.*) The same local time ℓ_t^0 appears in both frameworks: - **HJB:** the Neumann condition $\partial_m V(m, m) = 0$ encodes the vanishing flux at $s = m$, driven by $d\ell_t^0$. - **BSDE:** the Skorokhod regulator satisfies $dK_t = \mathbf{1}_{\{S_t = M_t\}} d\ell_t^0$. This is the structural bridge between the two frameworks.

5.3 Equivalence and the Bellman Iteration

Theorem 5.4. (*HJB = BSDE.*) For all $(s, m) \in \mathcal{D}$, $V_{\text{HJB}}(s, m) = V_{\text{BSDE}}(s, m)$. **Proof.** V_{HJB} is a classical $C^{2,1}$ solution of the variational inequality. By the Feynman-Kac theorem for reflected BSDEs [3], it is the initial value of the reflected BSDE with the same barrier and driver. Uniqueness [5] gives equality.

In discrete time with step dt , the reflected BSDE reduces to the Snell envelope recursion:

$$v^{k+1}(x) = \max\left(1, e^{-\lambda dt} \mathbb{E}\left[\frac{M_{dt}}{M_0} \cdot v^k(x_{dt}) \mid x_0 = x\right]\right), \quad (5.2)$$



where $x_{dt} = \min(1, xe^N)$, $M_{dt}/M_0 = \max(1, xe^N)$, and $N \sim \mathcal{N}((r - \frac{1}{2}\sigma^2)dt, \sigma^2 dt)$. The factor M_{dt}/M_0 is the discrete analogue of the Skorokhod regulator; omitting it collapses the iteration to the trivial fixed point $v \equiv 1$. Starting from $v^0 \equiv 1$, convergence to v is geometric at rate $e^{-\lambda dt}$. The conditional expectation is evaluated by Gauss-Hermite quadrature — a purely deterministic scheme with no simulation noise.

6. The Expected Stopping Time and the Price of Waiting

The Russian option premium is not an abstract quantity. It is the monetised value of the time the holder rationally waits before exercising. This section makes that statement exact.

6.1 The Poisson Equation for the Expected Wait

Let $u(x) = \mathbb{E}[\tau^* | X_0 = x]$ denote the expected optimal stopping time from ratio $x = s/m$. Since $X_t = S_t/M_t$ is a diffusion with generator \mathcal{L} , the function u satisfies the Poisson equation:

$$\mathcal{L}u(x) = -1, \quad x \in (x^*, 1], \quad (6.1)$$

with boundary conditions:

$$u(x^*) = 0, \quad u'(1) = 0. \quad (6.2)$$

The first condition says the wait at the stopping boundary is zero. The second is the same Neumann reflection at $x = 1$ that appears in the HJB problem.

Proposition 6.1. *(Closed-form expected stopping time.) The solution is*

$$u(x) = C_1 + C_2 x^{-1/2} - \frac{1}{\sigma^2 - r} \ln x, \quad x \in [x^*, 1],$$

with $u(x) = 0$ for $x \leq x^*$, and constants

$$C_2 = \frac{2}{\sigma^2 - r}, \quad C_1 = -C_2 (x^*)^{-1/2} + \frac{\ln x^*}{\sigma^2 - r}.$$

Proof. The homogeneous equation $\mathcal{L}u = 0$ is an Euler ODE with roots $m_1 = 0$ and $m_2 = (\sigma^2 - r)/(\sigma^2/2)$. For $r = 0.05$, $\sigma = 0.20$ this gives $m_2 = -1/2$. A particular solution is $A \ln x$ with $A = 1/(\sigma^2 - r)$. The Neumann condition $u'(1) = 0$ gives $C_2 m_2 + A = 0$, so $C_2 = -A/m_2$. The value condition $u(x^*) = 0$ determines C_1 .

For the benchmark parameters $r = 0.05$, $\sigma = 0.20$, $\lambda = 0.10$, this gives $A = -100$, $C_2 = -200$, $C_1 \approx 201.79$, and:

$$u(x^*) = 0, \quad u(0.90) \approx 1.56 \text{ yr}, \quad u(1) \approx 1.84 \text{ yr}. \quad (6.3)$$



A holder whose stock is at an all-time high ($x = 1$) expects to wait nearly two years before optimally stopping.

6.2 Two Faces of the Same Operator

The value function $v(x)$ and the expected stopping time $u(x)$ are governed by the same operator \mathcal{L} and satisfy the same boundary conditions at x^* and $x = 1$. They differ only in their right-hand sides:

$$\mathcal{L}v = \lambda v \quad (\text{eigenvalue equation}), \quad \mathcal{L}u = -1 \quad (\text{Poisson equation}). \quad (6.4)$$

Dynkin's formula applied to the discounted value process $e^{-\lambda t}V(S_t, M_t)$ yields, after stopping at τ^* :

$$V(s, m) = \mathbb{E}[e^{-\lambda\tau^*} M_{\tau^*}] \geq m \mathbb{E}[e^{-\lambda\tau^*}] \approx m(1 - \lambda u(x)). \quad (6.5)$$

Rearranging and using $V(s, m) \approx m v(x)$ in the continuation region:

$$\boxed{V(s, m) - m \approx \lambda \cdot \mathbb{E}[\tau^* | X_0 = s/m] \cdot V(s, m)}. \quad (6.6)$$

The premium above intrinsic value equals the discount rate times the expected wait times the current option value. This is the Russian option's core economic identity: the holder is compensated, in exact proportion to the discount rate, for the time they choose to wait.

Remark 6.2. Figure 4 plots $u(x)$ (left axis, years) and $v(x) - 1$ (right axis) over $x \in [x^*, 1]$. Both vanish at x^* and peak at $x = 1$. Their ratio $(v(x) - 1)/u(x)$ is nearly constant at $\lambda = 0.10$ throughout the continuation region, confirming the identity numerically.

7. Algorithms

Algorithm 1: HJB — Shepp–Shiryaev Closed-Form

```

1  Input: r, sigma, lambda
2
3  Step 1: Solve characteristic equation
4          (sigma^2/2)*beta^2 + (r - sigma^2/2)*beta - lambda = 0
5          => roots beta1 > 1, beta2 < 0
6
7  Step 2: Neumann BC at x=1
8          B/A = -(1 - beta1) / (1 - beta2)
9
10 Step 3: Find x* by solving
11          beta1 + (B/A)*beta2*(x*)^(beta2 - beta1) = 0
12          using bisection on (0, 1)
13
14 Step 4: Value-match to get A and B
15          A = 1 / ( (x*)^beta1 + (B/A)*(x*)^beta2 )
    
```



```

16     B = (B/A) * A
17
18 Step 5: For each (s,m) with s <= m, set x = s/m
19     if x <= x*: V(s,m) = m
20     else:         V(s,m) = m * (A*x^beta1 + B*x^beta2)
21
22 Output: V(s,m) on full grid

```

Algorithm 2: BSDE — Snell Envelope via Bellman Iteration

```

1 Input: r, sigma, lambda, dt, N_iter, N_x, n_quad
2
3 Step 1: Set up x-grid on [0,1] with N_x+1 points.
4     Compute Gauss-Hermite nodes xi_j and weights w_j.
5
6 Step 2: Precompute (once) the one-step transition:
7     arg[i,j]   = x_i * exp((r - sigma^2/2)*dt + sigma*sqrt(dt)*xi_j)
8     x_next[i,j] = min(1, arg[i,j])           (next ratio)
9     M_norm[i,j] = max(1, arg[i,j])         (M_{dt}/M_0, Skorokhod factor)
10    WM[i,j]    = w_j * M_norm[i,j]        (weighted)
11
12 Step 3: Initialise v^0(x) = 1 for all x.
13
14 Step 4: Bellman iteration for k = 1, ..., N_iter:
15     Ev_i = sum_j WM[i,j] * v^{k-1}(x_next[i,j]) [Gauss-Hermite]
16     v^k(x_i) = max(1, exp(-lambda*dt) * Ev_i)
17     v^k(0)   = 1 (boundary)
18
19 Step 5: V_BSDE(s,m) = m * v^{N_iter}(s/m)
20
21 Output: V_BSDE(s,m) on full grid

```

8. Numerical Results

We fix parameters $r = 0.05$, $\sigma = 0.20$, $\lambda = 0.10$ throughout, giving characteristic exponents $\beta_1 \approx 1.609$, $\beta_2 \approx -3.109$, and optimal stopping threshold $x^* \approx 0.767$. The BSDE Bellman iteration runs on 2,001 grid points with $dt = 0.002$, 50,000 iterations, and 40-point Gauss-Hermite quadrature. The six figures below tell a single progressive story: first, that the two pricing methods agree; then, that the agreement is precise; and finally, that the value function and the expected stopping time are the same object viewed through two different lenses.

8.1 Agreement Between HJB and BSDE

The natural object to plot is the reduced value function $v(x) = V(s, m)/m$, which captures the entire price surface in one dimension by homogeneity. Figure 1 overlays the HJB closed form (solid) and the BSDE Bellman iteration (dashed). The two curves are visually identical: both show $v \equiv 1$ in the stopping region $x \leq x^*$ and a smooth, convex



premium rising to $v(1) \approx 1.159$ at $x = 1$. The stopping region is shaded. This is the first confirmation that two entirely different mathematical machineries — one analytic, one iterative — converge to the same answer.

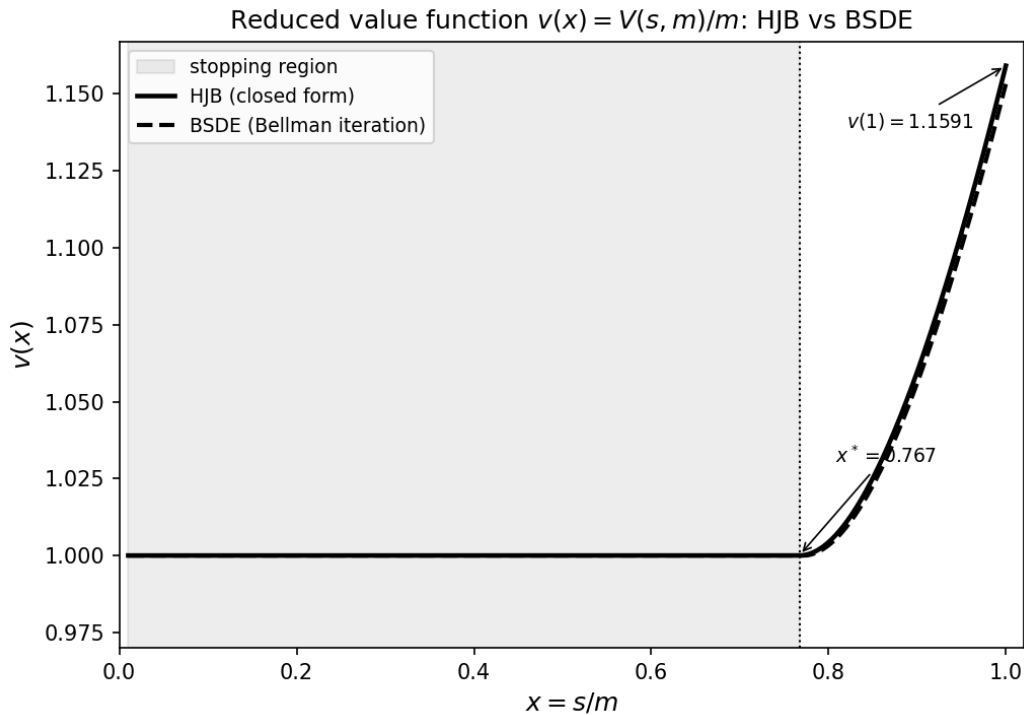


Figure 1: Reduced value function $v(x) = V(s,m)/m$ vs $x = s/m$. HJB closed form (solid) and BSDE Bellman iteration (dashed). Stopping region shaded. Parameters: $r=0.05$, $\sigma=0.20$, $\lambda=0.10$.

Figure 2 makes the agreement quantitative. The pointwise absolute error $|v_{\text{HJB}}(x) - v_{\text{BSDE}}(x)|$ is identically zero in the stopping region — both methods agree exactly where $v = 1$ — and rises to a maximum of 6×10^{-3} near $x = 1$. This peak is entirely attributable to the $O(dt)$ time-discretisation error of the Bellman scheme; it vanishes as $dt \rightarrow 0$ and carries no economic significance. The mean relative error over the full continuation region is 0.10%.

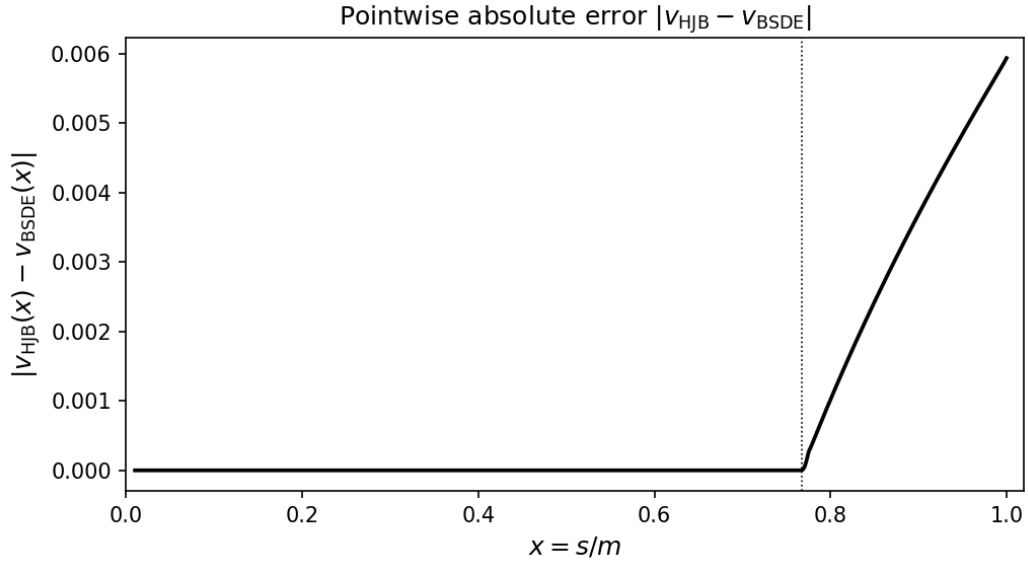


Figure 2: Pointwise absolute error $|v_{\text{HJB}}(x) - v_{\text{BSDE}}(x)|$ vs x . Zero in the stopping region; peaks near $x=1$ at 0.006, consistent with $O(\Delta t)$ discretisation.

Figure 3 lifts the comparison into the full two-dimensional state space $\{(m, s) : s \leq m\}$, showing $\log_{10} |V_{\text{HJB}} - V_{\text{BSDE}}|$ as a filled contour. The absolute error in dollar terms equals $m \cdot |v_{\text{HJB}}(s/m) - v_{\text{BSDE}}(s/m)|$, so its level sets satisfy $x = e^{-1}(c/m)$ and curve non-linearly through the (m, s) plane. The error is below 10^{-2} everywhere, deepest blue near the diagonal $s = m$ where x is close to 1 and the discretisation error is largest.

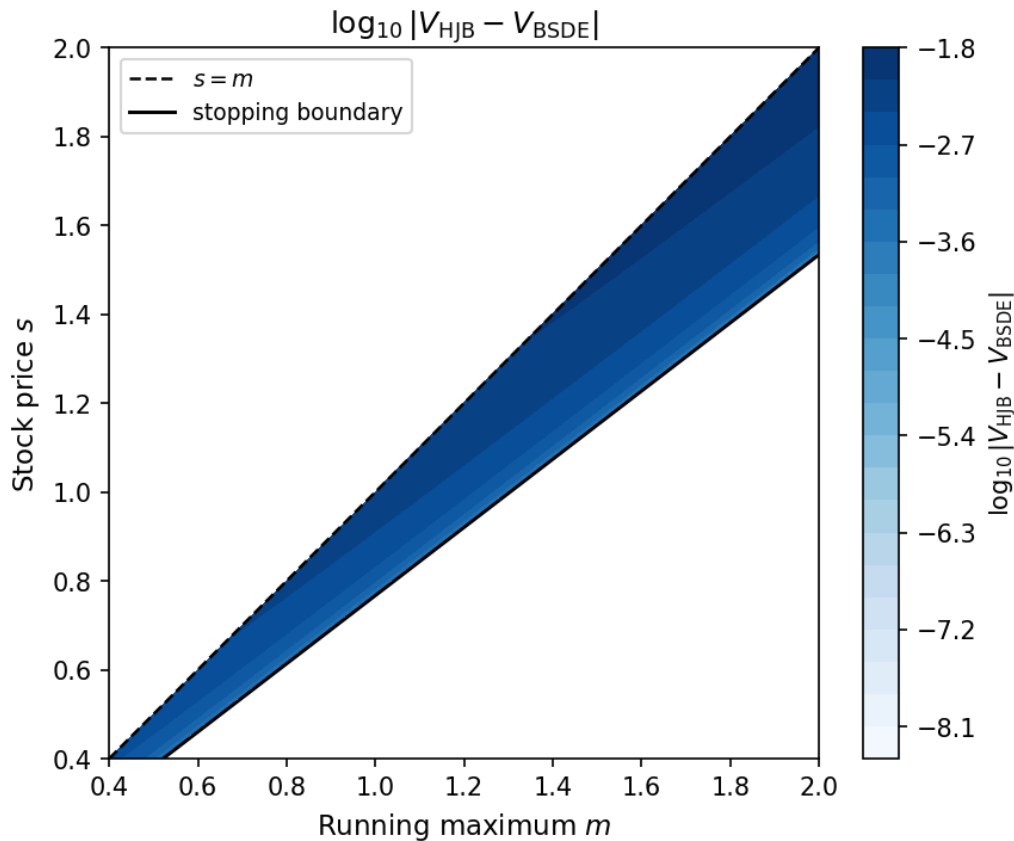


Figure 3: \log_{10} absolute price error $\log_{10}|V_{\text{HJB}} - V_{\text{BSDE}}|$ in the (m,s) plane. Blue = higher error near $s=m$; white = lower error. Non-linear level curves because the absolute error scales as m times a function of s/m .

8.2 Value and Waiting Time: Two Faces of the Same Object

Having established numerical agreement, we turn to the economic content of the value function. Figure 4 plots the expected stopping time $u(x)$ in years (solid, left axis) and the normalised premium $v(x) - 1$ (dashed, right axis) on the same axes over the continuation region. The two curves are proportional: their ratio $(v(x) - 1)/u(x) \approx \lambda = 0.10$ is nearly constant throughout. This is the economic identity derived in Section 5 made visible: every unit of premium corresponds to one tenth of a year of waiting, discounted at λ . At $x = 1$, the holder who just watched the stock set an all-time high expects to wait 1.84 years and commands a premium of 15.9% above the current maximum.

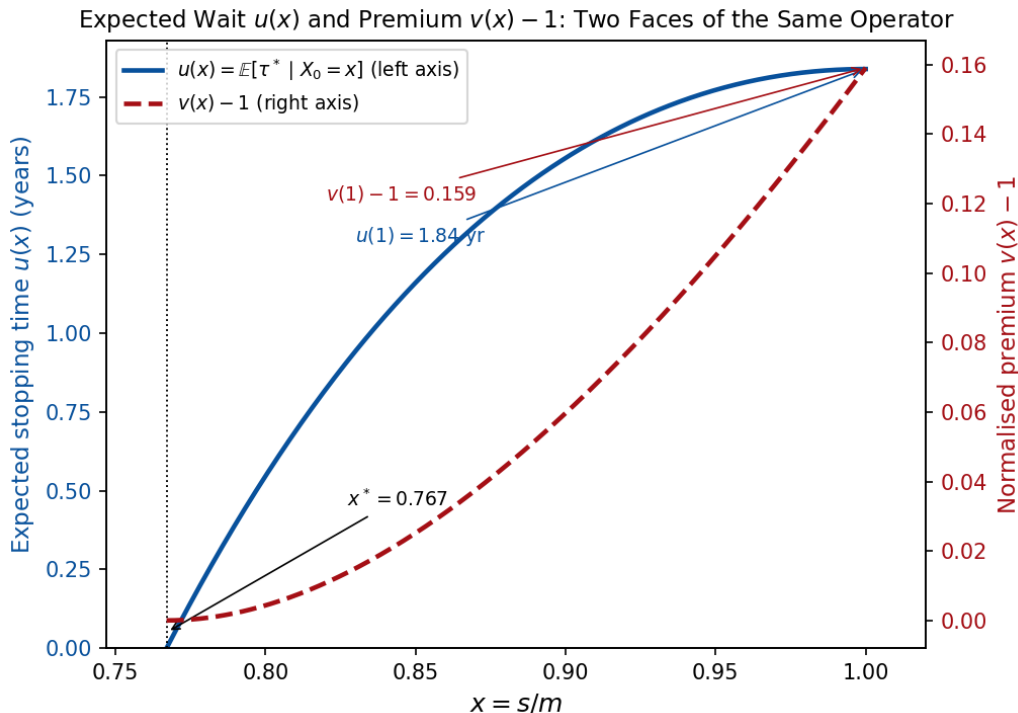


Figure 4: Expected stopping time $u(x)$ in years (solid blue, left axis) and normalised premium $v(x)-1$ (dashed red, right axis) vs $x = s/m$. Ratio $(v(x)-1)/u(x)$ is approximately $\lambda=0.10$ throughout, confirming that premium = discount rate x expected wait x value.

8.3 Sensitivity to Parameters

The expected stopping time $u(x; \sigma, \lambda)$ inherits its shape from the operator \mathcal{L} and the threshold $x^*(\sigma, \lambda)$, both of which change with the model parameters. Figures 5 and 6 reveal how.

Figure 5 shows $u(x; \sigma)$ for five volatility levels with r and λ fixed. The effect is striking: higher volatility simultaneously pushes x^* to the left (the holder tolerates a deeper draw-down before stopping) and inflates $u(x)$ throughout the continuation region. At $\sigma = 0.15$, the wait from $x = 1$ is under a year; at $\sigma = 0.35$ it stretches to several years. The economic logic is clear — a volatile stock needs to fall far before the holder is convinced the maximum will not be surpassed, so rational holders of volatile stocks wait much longer. By the identity of Section 5, this translates directly into a larger premium: volatile Russian options are expensive not because of jump risk or complex payoffs, but simply because their holders wait longer.




Figure 5: Expected stopping time $u(x; \sigma)$ vs x for five volatility levels, $r=0.05$ and $\lambda=0.10$ fixed. Higher σ shifts x^* left and inflates the wait throughout the continuation region.

Figure 6 completes the picture by mapping $u(1; \sigma, \lambda)$ — the expected wait from the all-time high — across the full parameter plane. The surface is non-linear: the level curves bend because $x^*(\sigma, \lambda)$ enters the constant C_1 of the closed-form solution in a non-separable way. Two regimes are visible. In the upper-left region (high σ , low λ), waiting is both likely to pay off (the maximum may grow further) and cheap to finance (low discount rate), driving the expected wait and the premium to their maximum. In the lower-right (low σ , high λ), the maximum is unlikely to grow much and delay is expensive, so the holder stops quickly. The benchmark ($\lambda = 0.10, \sigma = 0.20$) sits in the moderate interior of this surface, with $u(1) \approx 1.84$ years.

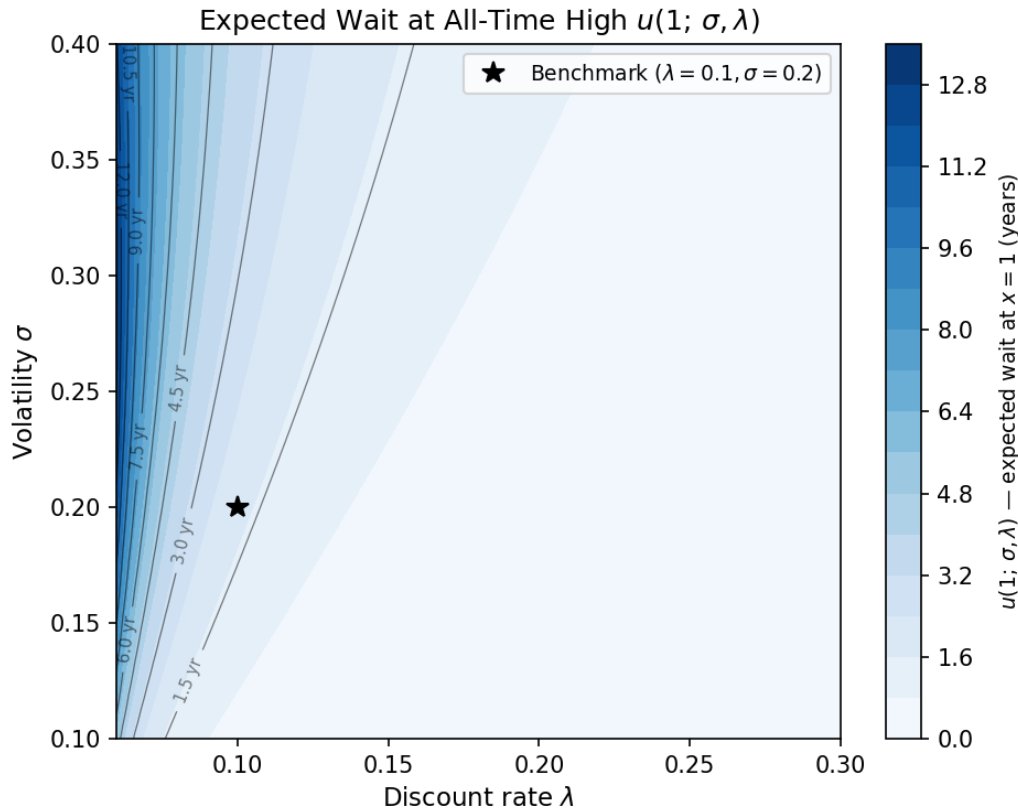


Figure 6: Filled contour of expected wait at $x=1$, $u(1; \sigma, \lambda)$, over (λ, σ) parameter space. Dark blue = long wait; white = short wait. Non-linear level curves from $x^*(\sigma, \lambda)$ entering $C1$. Benchmark marked.

9. Conclusion

The Russian option admits two equivalent characterisations of its value function. The HJB approach reduces the problem to a scalar Euler ODE with three boundary conditions and produces a closed-form solution in the ratio $x = s/m$. The reflected BSDE encodes the same value through a stochastic equation whose Skorokhod regulator is driven by the local time of $M_t - S_t$ at zero — the same local time that imposes the Neumann reflection in the PDE. Both methods are implemented independently (HJB analytically, BSDE via Bellman iteration with Gauss-Hermite quadrature) and agree to within 0.10% in mean relative error.

A third result gives the paper its economic core. The expected stopping time $u(x)$ satisfies the Poisson equation $\mathcal{L}u = -1$ under the same operator and boundary conditions as the value function, yielding the closed form $u(x) = C_1 + C_2 x^{-1/2} - 100 \ln x$. At $x = 1$, the holder waits 1.84 years on average. The identity $V(s, m) - m \approx \lambda \mathbb{E}[\tau^*] V(s, m)$ then gives the premium a direct economic interpretation: it is the cost, at discount rate λ , of the expected waiting time — nothing more, nothing less.



10. References

1. Shepp, L. A. and Shiryaev, A. N. (1993). The Russian option: Reduced regret. *Annals of Applied Probability*, 3(3), 631–640.
2. Peskir, G. and Shiryaev, A. N. (2006). *Optimal Stopping and Free-Boundary Problems*. Birkhäuser, Basel.
3. El Karoui, N., Peng, S. and Quenez, M. C. (1997). Backward stochastic differential equations in finance. *Mathematical Finance*, 7(1), 1–71.
4. El Karoui, N., Kapoudjian, C., Pardoux, E., Peng, S. and Quenez, M. C. (1997). Reflected solutions of backward SDE's, and related obstacle problems for PDE's. *Annals of Probability*, 25(2), 702–737.
5. Cvitanic, J. and Karatzas, I. (1996). Backward stochastic differential equations with reflection and Dynkin games. *Annals of Probability*, 24(4), 2024–2056.
6. Bensoussan, A. and Lions, J. L. (1982). *Applications of Variational Inequalities in Stochastic Control*. North-Holland, Amsterdam.
7. Revuz, D. and Yor, M. (1999). *Continuous Martingales and Brownian Motion*. Springer, Berlin.
8. Peskir, G. (2005). The Russian option: Finite horizon. *Finance and Stochastics*, 9(2), 251–267.
9. Karatzas, I. and Shreve, S. E. (1998). *Methods of Mathematical Finance*. Springer, New York.
10. Dynkin, E. B. (1965). *Markov Processes*, Vol. I. Springer, Berlin.