

TECHNICAL REPORT

Projected SOR, Penalty and Policy Iteration for Real-Options Variational Inequalities

A Comparison on Two-Factor Investment Problems

Supervised by Dr. T. Zamrik

08-MAY-2022

Penrose P3 tiling by Robinson-triangle inflation: rhombs of 36° and 72° , inflation factor $\varphi = \frac{1+\sqrt{5}}{2}$



TR-2026-08169923
zamrik.com



Zamrik Quant Research

The research arm of Zamrik Quant Lab

Copyright

© 2026 T. Zamrik. All rights reserved. Published by Zamrik Quant Research. No part of this document may be reproduced, stored, distributed or adapted, in whole or in part, by any means, without the prior written permission of the publisher, except for brief quotation with full citation. Requests for permission: info@zamrik.com.

Suggested Citation

Optimal Control & Decision Desk, Zamrik Quant Research (2022). 'Projected SOR, Penalty and Policy Iteration for Real-Options Variational Inequalities: A Comparison on Two-Factor Investment Problems'. *Zamrik Quant Research Technical Report TR-2026-08169923*, version 2 (13 October 2026). Available at <https://zamrik.com/research-items/tr-2026-08169923/>.

Earlier Version

First published on zamrik.com on 08-MAY-2022 as 'Projected SOR, Penalty and Policy Iteration for Real-Options Variational Inequalities' by T. Zamrik.

About Zamrik Quant Research

Zamrik Quant Research is the research arm of Zamrik Quant Lab. It works through five research desks, from market microstructure and derivatives to optimal control, stochastic analysis and applied mathematics, and publishes everything it finds in the open: research papers, notebooks that run on Kaggle, free tools that answer one practical question, and essays.

The five desks are Market Microstructure & Trading; Derivatives & Volatility; Optimal Control & Decision; Stochastic Analysis; and Applied Mathematics. Each paper, notebook, tool and essay belongs to exactly one desk. Zamrik Quant Research is run by a small, independent team, headed by Dr. T. Zamrik. We are self-funded: no sponsors, no advertisers, no outside capital.

Every publication passes through the same five functions before anyone outside sees it: research, data, software, verification and production. We never retrofit the mathematics to a conclusion chosen in advance. Every number is checked two independent ways. Every figure is re-run from its code. When we find an error, or replace a piece with a better one, we record it in the public corrections log at <https://zamrik.com/corrections/>.

Legal Notice

This report is research and education. Nothing in it is investment, financial, legal or engineering advice; nothing is a recommendation to buy or sell shares of stock, derivative contracts, bonds, or any other instrument or property; and nothing is a forecast. You are responsible for your own decisions. Figures and code are provided as they are. Care and independent checks reduce errors but cannot rule them out, so no warranty is given. The views expressed are those of the Optimal Control & Decision Desk.

TECHNICAL REPORT



Projected SOR, Penalty and Policy Iteration for Real-Options Variational Inequalities

A Comparison on Two-Factor Investment Problems

Supervised by Dr. T. Zamrik¹

08-MAY-2022

TR-2026-08169923

Abstract. An option to invest whose payoff depends on two random factors is valued by a variational inequality; after a monotone discretisation it is a linear complementarity problem with an M-matrix, which has one solution. This report compares projected SOR, the penalty method, policy iteration and, for an option that expires, Ikonen–Toivanen splitting, on three problems: the McDonald–Siegel option, checked against its closed form on the full two-factor grid; an option on a mean-reverting price, checked by simulating the rule its boundary defines; and an option with a finite life. All methods reach the same discrete solution, to 2.1×10^{-9} , and the boundary converges to the closed form $X^* = 2.8560$ at second order, to 2.4×10^{-4} on 512×512 cells. With a sparse direct solver on one thread, projected SOR was the fastest on the perpetual problems at the larger grids, 2.6 s against 26 s and 27 s at 512×512 . Policy iteration from the payoff moved the boundary about one cell per iteration, so its count grew with the grid, from 4 to 23. The penalty error fell as one over the penalty, which must be set by a rule. The nine-point stencil, never an M-matrix, converged and was the most accurate at strong correlation. For the expiring option, splitting with one factorisation matched the exact step solvers at a fraction of their time. The recommendations hold for a direct solver, one thread and problems of this size.

TR-2026-08169923 · Optimal Control & Decision Desk · Technical Report · Published
08-MAY-2022 · Updated 13-OCT-2026 · Reading time: about 36 min

1. Introduction

A firm that can build a project at a cost K and receive a project worth V holds an option: it may invest now, or wait. When both V and K move at random, this is an optimal stopping problem in two state variables. McDonald and Siegel [1] showed that, for geometric dynamics, the value is homogeneous of degree one, so the decision depends on V/K alone and the firm invests when the ratio first exceeds a threshold well above one. The same structure underlies the entry and exit decisions of Dixit [3] and the valuation of natural resources by Brennan and Schwartz [4]; Dixit and Pindyck [2] set out the theory.

As a partial differential problem, optimal stopping is a variational inequality: the value

¹Correspondence: Optimal Control & Decision Desk, Zamrik Quant Research, info@zamrik.com. Version 2 — re-issued by the Optimal Control & Decision Desk, 13 October 2026.



dominates the payoff, satisfies the pricing equation where it exceeds it, and the two conditions hold in complementary regions [6]. Jaillet, Lamberton and Lapeyre [7] gave the framework for American options, Peskir and Shiryaev [8] treat the free boundary, and Pham [9] places it within stochastic control. Without homogeneity the inequality must be solved on a two-dimensional grid, as a large complementarity problem whose method of solution decides the cost and how far the answer can be trusted.

Four methods dominate. Projected successive over-relaxation, introduced by Cryer [13] for quadratic programmes with bounds, sweeps the grid and clips each value at the obstacle; Brennan and Schwartz [11] had priced the American put by finite differences with the early-exercise condition applied node by node. It is short to write and needs no linear solver, but its theory covers symmetric matrices and its speed depends on a relaxation factor. The penalty method replaces the inequality by a nonlinear equation that penalises violations of the obstacle and solves it by Newton's method; Zvan, Forsyth and Vetzal [14] used it with stochastic volatility, and Forsyth and Vetzal [15] proved that under M-matrix conditions its Newton iteration is monotone and terminates. Policy iteration, Howard's algorithm [17], alternates between choosing on each node whether to invest and solving the system that choice defines; Bokanowski, Maroso and Zidani [18] proved its superlinear convergence, and Reisinger and Witte [20] proposed it as an easy way to price American options. For an option that expires, the splitting of Ikonen and Toivanen [22] divides each step into one linear solve and one pointwise update, so a single factorisation serves every step. Achdou and Pironneau [32] survey these methods. They are rarely compared on one discretisation, with one stopping test, against an exact answer, in two state variables.

The report makes four contributions. First, it states three investment problems with the facts the computation is checked against: a closed form and its one-dimensional reduction (Proposition 1, Corollary 1), the value of an output stream (Lemma 1) and the limit of a growing life (Proposition 2). Second, it gives the convergence theory in one place: when the seven-point stencil is an M-matrix (Proposition 3), uniqueness (Proposition 4), finite termination of policy iteration (Theorem 1), the penalty error with an explicit constant (Theorem 2), and policy iteration as a semismooth Newton method (Proposition 5). Third, it runs the four methods on one discrete problem per experiment, with one start, one stopping test and one thread, against exact answers where they exist and against refinement and simulation where they do not. Fourth, it reports where the usual advice fails, and states which method to use under which conditions (Remark 3) and what the comparison leaves out (Remark 4).

2. The Three Problems

Definition 1 (Investing when value and cost are both random). The value V of a completed project and the cost K of building it follow

$$dV_t = \mu_V V_t dt + \sigma_V V_t dW_t^1, \quad dK_t = \mu_K K_t dt + \sigma_K K_t dW_t^2,$$

with $d\langle W^1, W^2 \rangle_t = \rho dt$, $|\rho| < 1$, $\sigma_V, \sigma_K > 0$ and a discount rate $r > 0$; the shortfalls $\delta_V = r - \mu_V$ and $\delta_K = r - \mu_K$ are the rates at which waiting forgoes value and saves cost.



The option to invest is worth $F(V, K) = \sup_{\theta} \mathbb{E}[e^{-r\theta}(V_{\theta} - K_{\theta})^+]$ over stopping times θ , the payoff being zero on $\{\theta = \infty\}$. With

$$\begin{aligned} \mathcal{L} = & \frac{1}{2}\sigma_V^2 V^2 \partial_{VV} + \rho\sigma_V\sigma_K VK \partial_{VK} + \frac{1}{2}\sigma_K^2 K^2 \partial_{KK} \\ & + \mu_V V \partial_V + \mu_K K \partial_K, \end{aligned}$$

F solves $\min\{rF - \mathcal{L}F, F - (V - K)^+\} = 0$ on $V, K > 0$.

This is problem (a) [1]: the question is not whether V exceeds K but by how much it must, before waiting stops paying.

Proposition 1 (The McDonald–Siegel boundary). In Definition 1 let $\delta_V > 0$ and $\sigma^2 = \sigma_V^2 - 2\rho\sigma_V\sigma_K + \sigma_K^2 > 0$, let β be the root above one of

$$\frac{1}{2}\sigma^2\beta(\beta - 1) + (\delta_K - \delta_V)\beta - \delta_K = 0,$$

and $X^* = \beta/(\beta - 1)$. Then $F(V, K) = Kf(V/K)$ with $f(X) = (X^* - 1)(X/X^*)^{\beta}$ for $X < X^*$ and $f(X) = X - 1$ for $X \geq X^*$; it is optimal to invest the first time $V \geq X^*K$; and $X^* > \max(1, \delta_K/\delta_V)$.

Proof. The quadratic is $-\delta_V < 0$ at one, so it has one root above one. The dynamics and the payoff scale with (V, K) , so $F(cV, cK) = cF(V, K)$ and $F = Kf(V/K)$, $f(X) = F(X, 1)$. With $X = V/K$, $\partial_V F = f'$, $\partial_K F = f - Xf'$, $\partial_{VV} F = f''/K$, $\partial_{VK} F = -Xf''/K$, $\partial_{KK} F = X^2 f''/K$, and $\mu_V - \mu_K = \delta_K - \delta_V$, so

$$rF - \mathcal{L}F = K[\delta_K f - (\delta_K - \delta_V)Xf' - \frac{1}{2}\sigma^2 X^2 f''].$$

Below the boundary f combines X^b over the two roots b ; as $e^{-rt}V_t$ is a supermartingale, $f(X) \leq X$, which excludes the root below one. Value matching $A(X^*)^{\beta} = X^* - 1$ and smooth pasting $A\beta(X^*)^{\beta-1} = 1$ give X^* and f [1], [2]; f is convex and tangent to $X - 1$ there, so $f > (X - 1)^+$ below X^* . Above it, $rF - \mathcal{L}F = K(\delta_V X - \delta_K)$, and the quadratic gives $\frac{1}{2}\sigma^2\beta(\beta - 1) = \delta_V\beta - \delta_K(\beta - 1) > 0$, so $\delta_V X^* > \delta_K$. The verification argument [2] identifies F with the value. $\square$

The correlation enters only through σ^2 , the variance of $\ln(V/K)$; with Example 1's other parameters the boundary falls from 3.79 at $\rho = -0.95$ to 2.30 at $\rho = 0.95$.

Corollary 1 (One ratio carries the whole problem). Under Proposition 1, $F = Kf(V/K)$ solves the two-factor inequality of Definition 1 if and only if

$$\min\{\delta_K f - \frac{1}{2}\sigma^2 X^2 f'' - (\delta_K - \delta_V)Xf', f - (X - 1)^+\} = 0, \quad X > 0,$$

the inequality of a perpetual call struck at one, with volatility σ , yield δ_V and rate δ_K .

Proof. By the proof of Proposition 1, $rF - \mathcal{L}F$ and $F - (V - K)^+$ are $K > 0$ times the two expressions above. $\square$



The reduction is a second, independent discretisation of the same closed form, on a grid in $\ln X$ alone.

Example 1 (The base case in numbers). Problem (a) with $r = 0.05$, $\mu_V = 0.01$, $\sigma_V = 0.25$, $\mu_K = -0.02$, $\sigma_K = 0.10$ and $\rho = 0.3$: the value grows at 1 per cent a year and the cost falls at 2 per cent.

Parameters. $\delta_V = 0.04$, $\delta_K = 0.07$, $\sigma^2 = 0.0575$.

Computation. The quadratic $\frac{1}{2}(0.0575)\beta(\beta - 1) + (0.07 - 0.04)\beta - 0.07 = 0$ has the root $\beta = 1.5388$, so $X^* = 2.8560$, and $f(e^{-1}) = 0.0792$, $f(1) = 0.3692$, $f(e^{0.4}) = 0.6832$, $f(e^{0.8}) = 1.2644$. In the investment region $\delta_V X^* - \delta_K = 0.0442 > 0$. Numerically, the quadratic's residual is 1.4×10^{-17} , and $rF - \mathcal{L}F$ applied to $Kf(V/K)$ by finite differences in (V, K) vanishes to 7.2×10^{-9} relative.

Shows. The firm invests when the project is worth 2.856 times its cost, not 1; where investing now is worth nothing, $X = 1$, the option is worth 0.369 per unit of cost.

Definition 2 (A mean-reverting price and a falling cost). The project yields one unit of output a year for ever, sold at a price P with $d \ln P_t = \kappa(\ln \bar{p} - \ln P_t) dt + \sigma_P dW_t^1$, $\kappa > 0$; the cost K follows Definition 1, with $d\langle W^1, W^2 \rangle_t = \rho dt$. Investing at a stopping time θ pays $\Pi(P_\theta) - K_\theta$, where

$$\Pi(P) = \int_0^\infty e^{-rt} \mathbb{E}_P[P_t] dt,$$

and the option is worth $F(P, K) = \sup_\theta \mathbb{E}[e^{-r\theta}(\Pi(P_\theta) - K_\theta)^+]$. The base case has $\kappa = 0.3$, $\bar{p} = 1$, $\sigma_P = 0.3$, $r = 0.05$, $\mu_K = -0.02$, $\sigma_K = 0.10$ and $\rho = 0.2$.

This is problem (b), with the one-factor price of Schwartz [5], a case Dixit and Pindyck [2] discuss. Π is not linear in P , so there is no reduction and no closed form; in $x = \ln P$, $y = \ln K$ the diffusion is constant and the drift $\kappa(\ln \bar{p} - x)$ points back to $\ln \bar{p}$.

Lemma 1 (The value of the output stream). Under Definition 2, $\ln P_t$ is Gaussian with mean $m_t = e^{-\kappa t} \ln P + (1 - e^{-\kappa t}) \ln \bar{p}$ and variance $s_t^2 = \sigma_P^2(1 - e^{-2\kappa t})/(2\kappa)$, so $\mathbb{E}_P[P_t] = \exp(m_t + s_t^2/2)$; Π is finite, with $\Pi(P) \leq r^{-1} \max(P, \bar{p}) \exp(\sigma_P^2/(4\kappa))$, and continuous and strictly increasing in P .

Proof. The equation is linear in $\ln P_t$, so

$$\ln P_t = m_t + \sigma_P \int_0^t e^{-\kappa(t-s)} dW_s^1,$$

a Wiener integral with variance s_t^2 added to a constant, and the lognormal mean gives $\mathbb{E}_P[P_t]$. As m_t is a convex combination of $\ln P$ and $\ln \bar{p}$, $e^{m_t} \leq \max(P, \bar{p})$, and $s_t^2 \leq \sigma_P^2/(2\kappa)$, which gives the bound. The integrand increases strictly with P , at the rate $e^{-\kappa t}$ in $\ln P$, and is dominated on bounded sets of P , so Π is continuous and strictly increasing. $\square$

So the payoff of problem (b) increases in P and decreases in K . Since $e^{-rt}K_t$ is a supermartingale when $\delta_K > 0$, raising the starting cost by ΔK lowers F by at most ΔK ;



hence $F + K - \Pi$ is non-decreasing in K , and the investment region is $\{K \leq K^*(P)\}$, one boundary per price, which is what the computation locates and the simulation uses.

Definition 3 (An option to invest that expires). The option of Definition 1 may be exercised only until a date T . With τ the time left, its value $F(V, K, \tau)$ is the supremum of $\mathbb{E}[e^{-r\theta}(V_\theta - K_\theta)^+]$ over stopping times $\theta \leq \tau$; it solves

$$\min\{\partial_\tau F + rF - \mathcal{L}F, F - (V - K)^+\} = 0, \quad F(V, K, 0) = (V - K)^+,$$

and is homogeneous, $F = Kf(V/K, \tau)$, with investment optimal when $V \geq X^*(\tau)K$.

This is problem (c). In X it is an American call with yield δ_V and rate δ_K , so its boundary starts, as $\tau \rightarrow 0$, at $\max(1, \delta_K/\delta_V)$, which is 1.75 in Example 1. The payoff's kink at $X = 1$ then lies in the continuation region, which matters for the time stepping.

Proposition 2 (A longer life is worth more, up to the perpetual option). Under Definition 3 with $\delta_V > 0$, let F_T and $X^*(T)$ be the value and boundary with life T , and F, X^* those of Proposition 1. For $T \leq T'$, $F_T \leq F_{T'} \leq F$ and $X^*(T) \leq X^*(T') \leq X^*$; as $T \rightarrow \infty$, $F_T \rightarrow F$ and $X^*(T) \rightarrow X^*$.

Proof. Stopping times bounded by T are bounded by T' and admissible for the perpetual problem, which orders the values; the regions $\{F_T = (V - K)^+\} = \{V \geq X^*(T)K\}$ are then nested in reverse, which orders the boundaries. Fix a stopping time θ . As $T \rightarrow \infty$, $e^{-r(\theta \wedge T)}(V_{\theta \wedge T} - K_{\theta \wedge T})^+ \rightarrow e^{-r\theta}(V_\theta - K_\theta)^+$, which is zero on $\{\theta = \infty\}$ since $e^{-rt}V_t \rightarrow 0$. The payoffs are dominated by $\sup_t e^{-rt}V_t$, integrable because the supremum of $\sigma_V W_t^1 - (\delta_V + \frac{1}{2}\sigma_V^2)t$ is exponential with rate $(2\delta_V + \sigma_V^2)/\sigma_V^2 > 1$ [8]. Dominated convergence gives $\lim_T F_T \geq \mathbb{E}[e^{-r\theta}(V_\theta - K_\theta)^+]$ for every θ , so $\lim_T F_T \geq F$. Finally $X^*(T)$ increases to some $\bar{X} \leq X^*$; if $\bar{X} < X^*$, every $F_T(X, 1)$ with $\bar{X} < X < X^*$ is $X - 1$, so $f(X) = X - 1$ there, against Proposition 1. $\square$

Assumption 1 (Standing assumptions). (i) $r > 0$; $\sigma_V, \sigma_K, \sigma_P$ and κ are positive; $|\rho| < 1$. (ii) $\delta_V > 0$ and $\delta_K > 0$. (iii) Parameters are those of Example 1 and Definition 2 unless a study varies one.

Every statement below uses it: (i) gives the discrete matrices their diagonal dominance (Propositions 3 and 4); (ii) gives the closed form, the limit of the finite life and, in problem (b), the form of the investment region (Lemma 1); (iii) fixes the base cases on which the methods are compared.

3. Discretisation

Definition 4 (Grid, truncation and boundary conditions). Problems (a) and (c) use $x = \ln V, y = \ln K$ on $[-8.8, 4.0] \times [-2.56, 2.56]$ with $N \times N$ cells of sides $h_x, h_y, h_y/h_x = 0.4 = \sigma_K/\sigma_V$. The edge $x = -8.8$ ($V/K \leq e^{-6.24}$) carries $F = 0$, the edge $x = 4.0$ ($V/K \geq e^{1.44} = 4.22 > X^*$) carries $F = V - K$, and the edges $y = \pm 2.56$, which the boundary crosses, use $F(e^\ell V, e^\ell K) = e^\ell F(V, K)$, $\ell = 2h_x = 5h_y$: each edge node is $e^{\pm\ell}$



times the node two cells along and five inward. Problem (b) uses $x = \ln P$, $y = \ln K$ on $[-2, 2] \times [0, 5.5]$ with $N \times 4.125N$ cells, $h_y = h_x/3$; $F = \Pi(P) - K$ at $y = 0$, $F = 0$ at $y = 5.5$, and on the edges in x the equation without x -diffusion and cross term, its drift upwind inward. The reduction of Corollary 1 uses $z = \ln X$ on $[-10.4, 2.4]$ with N cells.

The homogeneity rows are exact for the discrete problem and use no closed-form value. Both grids have $a = 1$ in Proposition 3, so the seven-point stencil, used in every study but the correlation study, is monotone. Moving each edge outward changed the reference values by at most 4.5×10^{-9} in (a) and 1.3×10^{-7} in (b), relative.

The grid turns the inequality into a complementarity problem, and its form decides which theory applies.

Start. In the coordinates of Definition 4, with $F(V, K) = u(x, y)$, the generator of Definition 1 is

$$\mathcal{L}u = a_{11}\partial_{xx}u + a_{12}\partial_{xy}u + a_{22}\partial_{yy}u + b_1\partial_xu + b_2\partial_yu, \quad (3.1)$$

with $a_{11} = \frac{1}{2}\sigma_V^2$, $a_{12} = \rho\sigma_V\sigma_K$, $a_{22} = \frac{1}{2}\sigma_K^2$, $b_1 = \mu_V - \frac{1}{2}\sigma_V^2$, $b_2 = \mu_K - \frac{1}{2}\sigma_K^2$, and the inequality is $\min\{ru - \mathcal{L}u, u - g\} = 0$, $g = (e^x - e^y)^+$ [6], [32].

Steps. (1) At interior nodes a stencil replaces $\mathcal{L}$ by a row of L_h , giving $A = rI - L_h$ and $b = 0$; boundary rows are $u_i = g_i$ on the Dirichlet edges and $u_e - e^{\pm\ell}u_{e'} = 0$ on the homogeneity edges. Row by row,

$$\min(Au - b, u - g) = 0. \quad (3.2)$$

(2) As $\min(p, q) = 0$ exactly when $p, q \geq 0$ and $pq = 0$, this is the complementarity problem $u \geq g$, $Au \geq b$, $(Au - b)^\top(u - g) = 0$. (3) In problem (c) the θ -scheme gives at each step the same problem with $A_\theta = I + \theta\Delta t A$ and $b = (I - (1 - \theta)\Delta t A)u^n$; $\theta = 1$ is implicit Euler, $\frac{1}{2}$ Crank–Nicolson. (4) On the edge $y = 2.56$ the homogeneity rows have negative sums. The change of unknown $u = Ew$, $E = \text{diag}(e^{y_j})$, keeps every sign, turns them into $w_e - w_{e'} = 0$, and gives each interior row the sum $e^{-y_i}(Ae^y)_i$, which is δ_K up to $O(h_y^2)$, since $(r - \mathcal{L})e^y = \delta_K e^y$.

End. If the interior rows have non-positive off-diagonal entries, let z be one except slightly larger on the homogeneity edges; then $E^{-1}AEz > 0$, so $E^{-1}AE$, and with it A , is a nonsingular M-matrix by the criterion in the proof of Proposition 4. So is A_θ , and so is problem (b)'s matrix directly. Proposition 3 says when the interior rows qualify.

Proposition 3 (When a nine-point stencil stays monotone). Let $A = rI - L_h$ be assembled as in the derivation above with spacings h_x , h_y , central second differences on the axes, the drift central where the axis weight is at least $|b_k|/(2h)$ and upwind elsewhere, and one of three approximations of $a_{12}\partial_{xy}u$. Let $a = \sigma_V h_y / (\sigma_K h_x)$. (i) The central nine-point stencil S9, $(u_{1,1} - u_{1,-1} - u_{-1,1} + u_{-1,-1})/(4h_x h_y)$, gives an M-matrix only at $\rho = 0$. (ii) The seven-point stencil S7 oriented by the sign of ρ , for $\rho \geq 0$

$$\partial_{xy}u \approx \frac{2u_{0,0} + u_{1,1} + u_{-1,-1} - u_{1,0} - u_{-1,0} - u_{0,1} - u_{0,-1}}{2h_x h_y},$$



gives an M-matrix exactly when $|\rho| \leq \min(a, 1/a)$. (iii) The wide stencil W from Selling's decomposition gives an M-matrix at every $|\rho| < 1$. (iv) For S7 within its range and for W the scheme is monotone, stable and consistent, and converges to the viscosity solution whenever the limit problem has a strong comparison principle [26], [29].

Proof. The neighbour weights of L_h sum to zero, so if all are non-negative A is a Z-matrix strictly dominant by r , an M-matrix; a negative weight is a positive off-diagonal entry. (i) The corner weights are $\pm a_{12}/(4h_x h_y)$. (ii) For $\rho \geq 0$ the weights are $a_{11}/h_x^2 - |a_{12}|/(2h_x h_y)$ on $(\pm 1, 0)$, $a_{22}/h_y^2 - |a_{12}|/(2h_x h_y)$ on $(0, \pm 1)$ and $|a_{12}|/(2h_x h_y)$ on $\pm(1, 1)$; upwinding adds $|b_k|/h$ to one side only, so the drift never makes a weight negative, and the conditions $\sigma_V^2/h_x^2 \geq |\rho| \sigma_V \sigma_K/(h_x h_y)$ and $\sigma_K^2/h_y^2 \geq |\rho| \sigma_V \sigma_K/(h_x h_y)$ are $|\rho| \leq a$ and $|\rho| \leq 1/a$. (iii) Selling's decomposition writes the diffusion tensor in grid units as $\sum_k w_k \xi_k \xi_k^\top$ with $w_k \geq 0$ and integer ξ_k , a sum of second differences along ξ_k with non-negative weights [27], [28]. (iv) Stability is $\|A^{-1}\|_\infty \leq 1/r$ on dominant rows and consistency is Taylor's; the theorem is that of [26], and the weights divided by the diagonal are a Markov chain approximation [29]. $\square$

At $a = 1$ S7 is monotone at every correlation, and Selling's reduction returns S7 itself.

Proposition 4 (One discrete solution for every M-matrix). Let A be a nonsingular M-matrix and $F(v) = \min(Av - b, v - g)$ row by row. (i) $F(u) = 0$ has exactly one solution. (ii) Every matrix whose rows are rows of A or of the identity is a nonsingular M-matrix. (iii) If $F(v) \leq 0 \leq F(w)$, then $v \leq w$.

Proof. A Z-matrix is a nonsingular M-matrix exactly when it maps some $z > 0$ to a positive vector [12]; take such a z for A . (ii) Such a matrix B is a Z-matrix with $Bz > 0$, since its rows give $(Az)_i > 0$ or $z_i > 0$. (iii) Let B take in row i the row of A if $(Av - b)_i \leq (v - g)_i$ and the row of the identity otherwise, with $c_i = b_i$ or g_i . Then $Bv - c = F(v) \leq 0 \leq F(w) \leq Bw - c$, so $B(w - v) \geq 0$ and $w \geq v$ by (ii). (i) Uniqueness is (iii) applied both ways. An M-matrix is a P-matrix, and the complementarity problem with a P-matrix has a solution for every right-hand side [12]; Theorem 1 also constructs it. $\square$

Part (iii), a discrete comparison principle, carries Theorems 1 and 2. Uniqueness means any difference between the methods' answers on one grid is how far each iteration was taken.

Remark 1 (The wide stencil and what it costs). W buys monotonicity at any correlation with reach: on square cells, where $a = 2.5$ and S7 is monotone only for $|\rho| \leq 0.4$, Selling's reduction used directions reaching two cells for $0.45 \leq |\rho| \leq 0.8$ and three for $|\rho| \geq 0.85$ [27], [28].

The price has three parts: a wider coupling, so more fill in a factorisation and boundary layers as thick as the reach; second differences over steps $|\xi|h$, still second-order consistent but with an error constant growing like $|\xi|^2$; and, where the reduction drops an axis direction, a drift upwinded at first order. At $\rho = 0.95$ on square cells W's value error



was 3.9×10^{-2} , against 2.0×10^{-2} for S7 and 1.5×10^{-3} for S9. W is needed only when $|\rho| > \min(a, 1/a)$; with constant diffusion in log coordinates, as here, choosing h_y/h_x as the ratio of the volatilities makes $a = 1$ and removes the need.

4. Four Solvers

Algorithm 1 — Projected SOR

Inputs: A , b and g ; which rows are interior; the tolerance $\epsilon = 10^{-10} \max(1, \|g\|_\infty)$; a factor ω , or none. *Output:* u with $\|\min(Au - b, u - g)\|_\infty \leq \epsilon$, and the number of sweeps.

1. If no ω is given, bound the Jacobi radius by Gershgorin over the interior rows, $\rho_J \leq \max_i(1 - r/a_{ii})$ on problem (a), and set $\omega = 2/(1 + \sqrt{1 - \rho_J^2})$ (Young).
2. Set $u = g$.
3. For $i = 1, \dots, n$, using values already updated: on a boundary row solve its equation for u_i ; on an interior row set $z = (b_i - \sum_{j \neq i} a_{ij}u_j)/a_{ii}$ and $u_i = \max(g_i, u_i + \omega(z - u_i))$.
4. Stop if $\|\min(Au - b, u - g)\|_\infty \leq \epsilon$; otherwise return to step 3.

Cost. One pass over the non-zeros per sweep; the sweeps grew like N , so the time grows roughly as N^3 . Over-relaxing the homogeneity rows instead made it diverge from $\omega \approx 1.8$. The method is Cryer's [13]; on convergence see [12] and Remark 2.

Algorithm 2 — The penalty method with a Newton iteration

Inputs: A , b , g , the interior rows and ϵ as in Algorithm 1; a penalty λ , or none. *Output:* u with $\|\min(Au - b, u - g)\|_\infty \leq \epsilon$, and the number of Newton steps.

1. If no λ is given, set $\lambda = 10 \|(Ag - b)^+\|_\infty / \epsilon$, the norm over the interior rows.
2. Set $u = g$.
3. Let P be diagonal with $P_{ii} = 1$ on interior rows where $u_i < g_i$, else 0. Stop if P repeats the previous step's; report a cycle if it repeats an earlier one.
4. Factorise $A + \lambda P$ and solve $(A + \lambda P)u = b + \lambda Pg$.
5. Stop if $\|\min(Au - b, u - g)\|_\infty \leq \epsilon$; otherwise return to step 3.

Cost. One sparse LU factorisation per step. By Theorem 2 the rule bounds the obstacle's violation by $\epsilon/10$. The iteration is Forsyth and Vetzal's [15], used with stochastic volatility in [14] and for Hamilton-Jacobi-Bellman equations in [16].

Algorithm 3 — Policy iteration

Inputs: A , b , g and ϵ as in Algorithm 1. *Output:* u with $\min(Au - b, u - g) = 0$ to within ϵ , its exercise set, and the number of policies.

1. Set $u = g$.
2. Let $E = \{i : (u - g)_i < (Au - b)_i\}$, where investing is the better choice at u . Stop



if E repeats the previous set; report a cycle if it repeats an earlier one.

3. Form B and c : on rows in E the row of the identity and $c_i = g_i$; on the others the row of A and $c_i = b_i$.
4. Factorise B and solve $Bu = c$.
5. Stop if $\| \min(Au - b, u - g) \|_\infty \leq \epsilon$; otherwise return to step 2.

Cost. One sparse LU factorisation per policy; the sets shrink from $\{i : (Ag - b)_i > 0\}$ (Theorem 1). Howard [17] introduced the method; [18] proves its convergence for this class and [20] applies it to American options.

Algorithm 4 — Ikonen-Toivanen splitting for a time step

Inputs: A , g and the interior rows; T and a schedule of steps $(\Delta t, \theta)$: $\theta = 1$ (implicit Euler), $\theta = 1/2$ (Crank-Nicolson), or four implicit half steps and then Crank-Nicolson (Rannacher). *Output:* u at $\tau = T$, and a multiplier $\eta \geq 0$ with $(u - g)_i \eta_i = 0$ on interior rows.

1. Set $u = g$ and $\eta = 0$.
2. For each distinct $(\Delta t, \theta)$, factorise $I + \theta \Delta t A$ once.
3. At each step solve $(I + \theta \Delta t A) \tilde{u} = (I - (1 - \theta) \Delta t A) u + \Delta t \eta$, with $\eta = 0$ on boundary rows.
4. On interior rows set $u^{\text{new}} = \max(\tilde{u} - \Delta t \eta, g)$ and $\eta^{\text{new}} = \max(0, \eta + (g - \tilde{u})/\Delta t)$; recompute the boundary rows from their equations.
5. Replace u, η by the new values and take the next step.

Cost. Two triangular solves per step; one factorisation, two for Rannacher. Step 4 solves the split problem exactly, $(u^{\text{new}} - g)_i \eta_i^{\text{new}} = 0$, and the step's own complementarity problem only approximately. The splitting is Ikonen and Toivanen's [22], extended to stochastic volatility in [23].

Theorem 1 (Policy iteration stops in finitely many steps). Let A be a nonsingular M-matrix with n rows, and run Algorithm 3 without its tolerance test from $u^0 = g$, with sets $E^0, E^1, \dots$, where u^{k+1} solves the system of E^k . Then (i) $g \leq u^1 \leq u^2 \leq \dots \leq u$, the solution of Proposition 4; (ii) $E^0 \supseteq E^1 \supseteq \dots$, strictly until a set repeats; (iii) the algorithm stops at u after at most $|E^0| + 1 \leq n + 1$ linear solves. More generally, Howard's algorithm with M-matrix policies converges superlinearly, and for obstacle problems in finitely many iterations, bounded by the number of unknowns [18].

Proof. Let $F(v) = \min(Av - b, v - g)$ and $B_k u^{k+1} = c_k$ the system of E^k ; by the choice of E^k , $B_k u^k - c_k = F(u^k)$. Now $F(u^0) = \min(Ag - b, 0) \leq 0$, and $F(u^k) \leq B_{k-1} u^k - c_{k-1} = 0$ for $k \geq 1$. So $B_k(u^{k+1} - u^k) = -F(u^k) \geq 0$ and $u^{k+1} \geq u^k$ by Proposition 4(ii), while $F(u^k) \leq 0 = F(u)$ gives $u^k \leq u$ by Proposition 4(iii): (i). If $i \notin E^k$, then $(Au^{k+1} - b)_i = 0 \leq (u^{k+1} - g)_i$, so $i \notin E^{k+1}$: (ii). If $E^{k+1} = E^k$, on its rows $u_i^{k+1} = g_i$ and $(Au^{k+1} - b)_i > 0$, and on the others $(Au^{k+1} - b)_i = 0 \leq (u^{k+1} - g)_i$, so $F(u^{k+1}) = 0$. A strictly decreasing chain of subsets of E^0 has at most $|E^0| + 1$ members: (iii). $\square$



Each policy can only release nodes, so the count measures how far the boundary must retreat from the first set, which the bound $n + 1$ does not.

Theorem 2 (The penalised solution and its error). Let A be a nonsingular M-matrix with n rows, $\lambda > 0$ and u the solution of Proposition 4. (i) $Av - b - \lambda(g - v)^+ = 0$ has exactly one solution u_λ . (ii) The Newton iteration of Algorithm 2 without its tolerance test increases from its first iterate on, its sets shrink, and it stops at u_λ after at most $n + 2$ linear solves. (iii) $0 \leq u - u_\lambda \leq C/\lambda$; if the rows of A have non-negative sums,

$$0 \leq u - u_\lambda \leq \|(Ag - b)^+\|_\infty / \lambda \quad \text{and} \quad (g - u_\lambda)^+ \leq \|(Ag - b)^+\|_\infty / \lambda.$$

Proof. As $-\lambda(g - v)^+ = \min_{p \in \{0,1\}} \lambda p(v - g)$, the equation is $\min_P [(A + \lambda P)v - b - \lambda P g] = 0$ over diagonal 0-1 matrices P , whose matrices satisfy $(A + \lambda P)z \geq Az > 0$ for the z of Proposition 4 and are M-matrices. Algorithm 2 takes $P_{ii} = 1$ where $v_i < g_i$, attaining the minimum: it is Howard's algorithm, and from $k = 1$ the proof of Theorem 1 applies, giving (ii) and existence. If v, v' both solve, $(A + \lambda D)(v - v') = 0$ with D diagonal in $[0, 1]$, so $v = v'$. For (iii), $Au_\lambda - b = \lambda(g - u_\lambda)^+ \geq 0$, so $\min(Au_\lambda - b, u_\lambda - g) \leq 0$ and $u_\lambda \leq u$. With non-negative row sums, at the i maximising $(g - u_\lambda)_i > 0$ the non-positive off-diagonal entries and the row sum give $(A(u_\lambda - g))_i \leq 0$, so $\lambda(g - u_\lambda)_i \leq (Ag - b)_i$; then $v = u_\lambda + \lambda^{-1}\|(Ag - b)^+\|_\infty \mathbf{1}$ has $v \geq g$ and $Av \geq b$, so $v \geq u$ by Proposition 4(iii). In general apply this to $Z^{-1}AZ$, $Z = \text{diag}(z)$ [15], [16]. $\square$

The rule of Algorithm 2 is this bound read backwards. The bound improves without limit in λ ; rounding sets the limit.

Policy iteration looks like a heuristic, a choice alternating with a solve, but it is Newton's method.

Start. The discrete problem is $F(u) = 0$ with $F_i(u) = \min((Au - b)_i, u_i - g_i)$, a piecewise linear, hence semismooth, map. A generalised derivative $J(u)$ takes as row i the row of A where $(Au - b)_i \leq (u - g)_i$ and the row of the identity otherwise [19].

Steps. With $E^k = \{i : (u^k - g)_i < (Au^k - b)_i\}$, the Newton step $J(u^k)(u^{k+1} - u^k) = -F(u^k)$ reads

$$u_i^{k+1} - u_i^k = -(u_i^k - g_i) \quad (i \in E^k), \quad (A(u^{k+1} - u^k))_i = -(Au^k - b)_i \quad (i \notin E^k), \quad (4.1)$$

that is, $u_i^{k+1} = g_i$ on E^k and $(Au^{k+1})_i = b_i$ elsewhere. In the primal-dual active set strategy [19] the multiplier of $u \geq g$ is $\mu = Au - b$, and the next active set is $\{i : \mu_i^k - c(u^k - g)_i > 0\}$; with $c = 1$ this is E^k .

End. The semismooth Newton step, the active set step with $c = 1$ and the step of Algorithm 3 solve the same system, so each inherits the others' theory: finite termination from Theorem 1, local superlinear convergence from [19].

Proposition 5 (Policy iteration is a semismooth Newton method, and the penalty iteration approaches it). Let A be a nonsingular M-matrix. (i) From the same iterate,



Algorithm 3, the semismooth Newton method and the primal–dual active set strategy with $c = 1$ give the same next iterate. (ii) At every iterate u^k , $k \geq 1$, of Algorithm 2 the penalised set $\{i : u_i^k < g_i\}$ is the exercise set Algorithm 3 would choose at u^k . (iii) For a fixed set P , the solution of $(A + \lambda P)v = b + \lambda Pg$ converges at the rate $1/\lambda$ to that of Algorithm 3’s system for the same set.

Proof. is the derivation above [19], [18]. (ii) u^k solves $(A + \lambda P^{k-1})u^k = b + \lambda P^{k-1}g$, so $(Au^k - b)_i = \lambda(g - u^k)_i$ on the rows of P^{k-1} and 0 elsewhere; in either case $(u^k - g)_i < (Au^k - b)_i$ holds exactly when $u_i^k < g_i$. (iii) Dividing the rows of P by λ gives $M_\lambda v = c_\lambda$ with $M_\lambda \rightarrow B$, $c_\lambda \rightarrow c$, the system of Algorithm 3, and B is nonsingular by Proposition 4(ii), so $v - B^{-1}c = O(1/\lambda)$ [15], [21]. $\square$

From $u^0 = g$ the penalty method’s first set is empty, so its first solve ignores the obstacle, a solve policy iteration does not make. Huang, Forsyth and Labahn [21] use the same identification for Hamilton–Jacobi–Bellman equations.

Remark 2 (What each method needs in order to converge). Each guarantee rests on the M-matrix, and each method has its own margin beyond it.

For projected Gauss–Seidel, $\omega = 1$, the M-matrix suffices: each update is non-decreasing in the other values, so from the payoff the iterates rise to a fixed point no larger than the solution, which by Proposition 4 is the solution. Over-relaxation breaks this monotonicity, and Cryer’s theorem for $0 < \omega < 2$ [13] needs a symmetric positive definite matrix, which drift and a cross derivative rule out; the measured runs failed from $\omega = 1.85$ at $N = 64$. Policy iteration needs every policy matrix to be an M-matrix, which Proposition 4 supplies [18]; otherwise a policy may recur, hence the cycle check of Algorithm 3. The penalty method needs the same [15], and a ceiling on λ : once λ times the machine precision times $\|g\|_\infty$ nears the test, a node just above the obstacle is rounded onto it and stays penalised. Splitting iterates nothing; its accuracy is measured.

5. How the Comparison Is Run

The comparison holds everything fixed except the method. Each experiment assembles one matrix, one right-hand side and one obstacle, and every method receives the same three. Every method starts from the payoff, $u = g$; in time stepping, the three iterative methods start each step from the previous solution. All stop on one test, $\|\min(Au - b, u - g)\|_\infty \leq 10^{-10} \max(1, \|g\|_\infty)$, checked after each sweep, Newton step or policy; on problem (a) it is 5.45×10^{-9} . All run on one thread, and a time is the median of three runs (one for problem (b) on its finest grid).

The computation is written in Python with NumPy, SciPy and Numba. Projected SOR’s sweeps are compiled before any run is timed. The penalty method, policy iteration and splitting share one sparse direct LU factorisation with a fill-reducing ordering, and splitting reuses one factorisation for all steps of a size. Times are comparable as far as the linear algebra is. All 104 reported solutions meet the test; the study took 862 s.



6. The McDonald-Siegel Problem on a Two-Factor Grid

Script. Maps the policy-iteration solution of problem (a) on the 256×256 grid.

Shows. Proposition 1 on the full two-factor grid.




Figure 1: The computed value of the option to invest $F(V, K)$ for problem (a) on the 256×256 grid by policy iteration at $\rho = 0.3$, with value contours, the closed-form boundary $V = 2.8560 K$ dashed and the computed boundary on sampled rows of fixed K as dots.

The firm waits above the dashed ray and invests below it. The dots lie on the ray along its length, as homogeneity requires; the boundary is 1.2×10^{-3} below the closed form, with a spread of 3.6×10^{-3} between rows.

Accuracy against the closed form. *Method.* Algorithms 1, 2 and 3 on problem (a), seven-point stencil. On each row of fixed K the boundary is the vertex of the parabola through the last three continuation nodes, the shape smooth pasting gives; X^* is its mean over the central half of the rows. Values are compared at $X = e^{-1}, 1, e^{0.4}, e^{0.8}$ on the row $K = 1$.

Parameters. Example 1; $N = 64$ to 512, and 1024 for projected SOR.

Result. The boundary's relative error is -2.14×10^{-2} , -6.23×10^{-3} , -1.22×10^{-3} and -2.44×10^{-4} , orders 1.78, 2.35, 2.33; at $N = 512$ it is 2.85531 against 2.85601. The value error is 1.27×10^{-3} , 3.18×10^{-4} , 8.31×10^{-5} and 2.12×10^{-5} , orders 1.99, 1.94, 1.97. At $N = 1024$: -1.13×10^{-4} and 5.45×10^{-6} .

Error. Second order in the value, and in the boundary to within the rows' scatter: at $N = 1024$ its order falls to 1.11, where its error is below the spread between rows, 3.0×10^{-4} . The methods' boundaries agree to ten significant figures and their solutions to 2.1×10^{-9} , so the error is the grid's. This verifies Proposition 1 on the two-factor grid.

The one-dimensional reduction as a second check. *Method.* The inequality of Corollary 1 on the grid in $z = \ln X$ of Definition 4, by Algorithms 1, 2 and 3, compared



with Proposition 1 at the same four points.

Parameters. Example 1 through $\sigma^2 = 0.0575$, $\delta_V = 0.04$, $\delta_K = 0.07$; $N = 64$ to 8192.

Result. At $N = 8192$ the boundary's relative error is -5.9×10^{-7} and the value's 6.5×10^{-7} , with orders 2.10, 2.27, 2.95 and 2.01, 2.05, 2.21 over the last three refinements. Policy iteration took 3, 6, 10, 20, 40, 79, 157 and 314 policies; the three methods agree on every grid.

Error. The reduction confirms the closed form, and Corollary 1, to 6.5×10^{-7} . Its boundary error is not monotone on coarse grids, 5.8×10^{-4} at $N = 128$ and 1.5×10^{-3} at 256, because the vertex estimate depends on where the boundary falls between nodes. At equal N its value error exceeds the two-factor grid's, 2.0×10^{-4} against 2.1×10^{-5} at $N = 512$: a different constant, the same order and the same limit.

Source. The accuracy study of problem (a): seven-point stencil on an aspect-matched grid, $\rho = 0.3$, every method stopped by the common test; the boundary X_h^* is the mean over rows of fixed K of the vertex estimate where the solution first meets the payoff; values at $X = e^{-1}, 1, e^{0.4}, e^{0.8}$; observed orders from successive halvings of the cell. At $N = 1024$ only projected SOR was run.

Columns. Cells per side; the methods that produced the row (they agree to 2×10^{-9}); relative error of the boundary and its observed order; largest relative value error and its observed order.

Table 1: Boundary and value errors of problem (a) against the closed form as the grid is refined.

N	methods	$X_h^*/X^* - 1$	order	largest value error	order
64	all three	-2.14×10^{-2}		1.27×10^{-3}	
128	all three	-6.23×10^{-3}	1.78	3.18×10^{-4}	1.99
256	all three	-1.22×10^{-3}	2.35	8.31×10^{-5}	1.94
512	all three	-2.44×10^{-4}	2.33	2.12×10^{-5}	1.97
1024	projected SOR	-1.13×10^{-4}	1.11	5.45×10^{-6}	1.96

The value errors fall at second order throughout. The boundary errors fall unevenly, at observed orders between 1.1 and 2.4, because the boundary is read from where the solution meets the payoff, which moves through the grid in discrete steps; the one-dimensional reduction, solved in X alone, reaches a boundary error of -5.9×10^{-7} and a value error of 6.5×10^{-7} at 8192 cells.

Source. The cost study of problem (a): seven-point stencil, $\rho = 0.3$, from the payoff to the common test; time of one solve, median of three runs, one thread.

Columns. Cells per side; unknowns; Young's factor; each method's iterations and time.



Table 2: Iterations and time to the common stopping test by solver and grid, problem (a).

N	unknowns	Young's ω	projected SOR	penalty method	policy iteration
64	4,225	1.679	65 sweeps, 0.0047 s	6 steps, 0.030 s	4 policies, 0.019 s
128	16,641	1.824	132 sweeps, 0.040 s	9 steps, 0.25 s	7 policies, 0.18 s
256	66,049	1.908	267 sweeps, 0.31 s	14 steps, 2.41 s	12 policies, 1.70 s
512	263,169	1.953	545 sweeps, 2.55 s	25 steps, 27.1 s	23 policies, 25.7 s
1024	1,050,625	1.976	1117 sweeps, 21.2 s	not run	not run

Projected SOR at $N = 1024$ was faster than either factorising method at $N = 512$.

Cost on equal footing. *Method.* Algorithms 1, 2 and 3 on the problems of the accuracy study, from the payoff, to the common test, timed as in the setup.

Parameters. Problem (a); 4,225 to 263,169 unknowns, 1,050,625 for projected SOR.

Result. Table 2. Projected SOR is the fastest at every grid, 0.0047 s to 2.55 s against 0.019 s to 25.7 s for policy iteration and 0.030 s to 27.1 s for the penalty method. Its sweeps double with N , as the Gershgorin bound $1 - r/a_{ii}$ rises from 0.9815 to 0.99993; the discount keeps it below one, and a sweep is cheap. The others pay a factorisation per iteration.

Error. The solutions differ by at most 2.1×10^{-9} , inside the test. Run to 5.5×10^{-13} at $N = 128$, projected SOR came within 1.1×10^{-13} of policy iteration: the gap at the common test is its last sweep's residual. The penalty method made exactly two more solves than policy iteration on every grid, as Proposition 5 leads one to expect.

How many policies policy iteration visits. *Method.* Algorithm 3 from the payoff, counting linear solves to the common test on problems (a) and (b) and the reduction, against Theorem 1.

Parameters. (a): $N = 64$ to 512; (b): $N = 32$ to 256 cells in $\ln P$; reduction: $N = 64$ to 8192.

Result. (a): 4, 7, 12 and 23 policies for 4,225 to 263,169 unknowns. (b): 8, 13, 26 and 50 for 4,389 to 271,649. Reduction: 3, 6, 10, 20, 40, 79, 157 and 314 for 65 to 8,193. Each count about doubles with N .

Error. Every count lies far inside the bound, 23 against 263,169. The proof explains the growth: the first set is where $Ag - b > 0$, which in problem (a) begins near $X = \delta_K/\delta_V = 1.75$, and each policy releases about one layer of cells until the set retreats to $X^* = 2.856$, so halving the cell size doubles the count. The final residuals, 1.5×10^{-14} to 2.6×10^{-12} , are those of an exact solve. This verifies Theorem 1: finite, but not flat.

Script. Plots the relaxation study: projected SOR sweeps against the relaxation factor on two grids.



Shows. Where Young's rule puts projected SOR relative to the best factor, and how sharply the cost rises past it.




Figure 2: Projected SOR sweeps to the common stopping test against the relaxation factor ω on problem (a) at $\rho = 0.3$, for $N = 64$ and $N = 128$, with Young's factor starred on each curve and runs that did not converge within 20,000 sweeps crossed.

Both curves fall slowly to a sharp minimum and then rise steeply. The best factor moves toward 2 as the grid is refined, from 1.600 (48 sweeps) at $N = 64$ to 1.775 (97 sweeps) at $N = 128$; Young's factor, 1.679 and 1.824, sits just past the minimum and costs about 35 per cent more sweeps, 65 and 132. Past the minimum the penalty for over-relaxing is far larger than for under-relaxing, and from $\omega = 1.85$ at $N = 64$ and $\omega = 1.925$ at $N = 128$ the runs do not converge within 20,000 sweeps, so a factor estimated slightly low is the safer error.

Projected SOR and its relaxation factor. *Method.* Algorithm 1 on problem (a) for $\omega = 1, 1.025, \dots, 1.95$ from the payoff, capped at 20,000 sweeps; and once to a test of 5.5×10^{-13} .

Parameters. Example 1; $N = 64$ and 128.

Result. At $N = 64$ the fewest sweeps, 48, came at $\omega = 1.600$; Young's 1.679 took 65 and $\omega = 1$ took 262; the 34 factors up to 1.825 converged and none from 1.85. At $N = 128$: 97 at 1.775, 132 at Young's 1.824, 1090 at $\omega = 1$, no convergence at 1.925 and 1.95. Young's rule costs about 35 per cent more sweeps than the best factor.

Error. Every converged run met the test, and to 5.5×10^{-13} projected SOR met policy iteration's solution within 1.1×10^{-13} , as Proposition 4's uniqueness requires. The failures do not contradict it: for $\omega > 1$ on a non-symmetric matrix no theorem says the iteration finds the solution (Remark 2).

The penalty error against the penalty. *Method.* Algorithm 2 with $\lambda = 10^2, \dots, 10^{10}$, run until its set repeats, against policy iteration's exact solution; then time stepping of



problem (c) with a fixed $\lambda = 10^{10}$ and with the rule.

Parameters. Problem (a), $N = 128$.

Result. The sup-norm error falls from 1.9×10^{-2} at $\lambda = 10^2$ to 2.0×10^{-10} at 10^{10} , slope -0.9987 over 10^2 to 10^9 , above the rounding floor of 1.2×10^{-11} ; every λ took nine Newton steps. The rule gives 3.3×10^9 to 3.9×10^9 on problem (a).

Error. The slope agrees with Theorem 2's -1 to three digits, and the sets, not λ , decide the step count. In time stepping a fixed $\lambda = 10^{10}$ reaches a rounding floor of about 10^{-5} , below which a released node stays penalised: up to 29 of 130 steps failed the test, with residuals up to 3×10^{-5} and errors up to 2.6×10^{-5} . With the rule none failed, the worst residual being 5.26×10^{-9} against 5.45×10^{-9} . This verifies Theorem 2 and its rule.

7. Correlation and Monotonicity

Script. Plots the correlation study for each stencil and grid.

Shows. Proposition 3's condition, governing convergence but not accuracy.




Figure 3: Left: largest relative value error at $X = e^{-1}, 1, e^{0.4}, e^{0.8}$ against the correlation ρ for the nine-point, seven-point and wide stencils on the aspect-matched grid ($a = 1$) and on square cells ($a = 2.5$); right: projected Gauss-Seidel sweeps at $\omega = 1$; filled markers are M-matrices, and shading marks $|\rho| > 0.4$.

S9 is hollow except at $\rho = 0$, yet the most accurate at strong correlation. On square cells S7 turns hollow in the shaded band, but its error jumps earlier, between $|\rho| = 0.3$ and 0.35 , where its drift is first upwinded. Every run converged.

The sweep counts on the right fall in straight lines with $|\rho|$ and meet in a corner at $\rho = 0$. The stencils depend on $|\rho|$: as the correlation grows, part of the diffusion leaves the two axes for a diagonal direction, $(1, 1)$ when $\rho > 0$ and $(1, -1)$ when $\rho < 0$, and the diagonal entry of each row falls linearly with $|\rho|$; projected Gauss-Seidel needs roughly that entry



over the discount in sweeps. On the aspect-matched grid, where the wide and seven-point stencils coincide, the count falls from 1204 at $\rho = 0$ to 602 at $\rho = -0.95$ and 715 at $\rho = 0.95$, near the factor $(2 - |\rho|)/2 = 0.525$ this rough argument gives. The two signs differ by about a sixth; the drift, which in log coordinates has two components of the same sign, is the likely cause, though the study does not isolate it. On square cells the wide stencil adds corners where Selling's decomposition needs longer steps, its reach growing from one cell at $|\rho| \leq 0.4$ to three at $|\rho| = 0.95$, and its count falls much faster, to 115 and 144: the longer steps act like a coarser grid, which converges quickly to a less accurate answer, as the left panel shows.

Where the standard stencil fails. *Method.* Problem (a) at 39 correlations from -0.95 to 0.95 with S9, S7 and W on two grids: positive off-diagonal entries counted against Proposition 3, projected Gauss–Seidel ($\omega = 1$, cap 20,000) and policy iteration run, values compared with the closed form at each ρ .

Parameters. Matched grid $h_x = 0.1$, $h_y = 0.04$, $a = 1$; square cells $h_x = h_y = 0.1$, $a = 2.5$, $\min(a, 1/a) = 0.4$.

Result. The count agrees with Proposition 3 at every correlation, stencil and grid: S9 has 34,290 positive entries (matched) and 13,770 (square) at every $\rho \neq 0$; S7 has none on the matched grid and none on square cells exactly when $|\rho| \leq 0.4$; W has none. Projected Gauss–Seidel converged everywhere. Policy iteration did too, but on S9 took 109 to 113 policies at every $\rho \geq 0.05$ on square cells and at $\rho = 0.8, 0.85, 0.95$ on the matched grid, against 4 to 9 on S7 and W. At $\rho = 0.95$ the value error is 2.3×10^{-3} (S9) and 6.5×10^{-3} (S7) on the matched grid, and 1.5×10^{-3} , 2.0×10^{-2} and 3.9×10^{-2} (S9, S7, W) on square cells.

Error. Monotonicity guaranteed convergence, not accuracy. Part of the loss is the upwinded drift: on square cells S7's error rises from 9.9×10^{-4} at $\rho = 0.3$ to 8.2×10^{-3} at 0.35 , where it starts to upwind, before it stops being an M-matrix at 0.4 . S9 upwinds nowhere; its missing guarantee showed only in the policy count. This verifies Proposition 3 exactly.

8. A Mean-Reverting Price

Script. Maps the policy-iteration value of problem (b), with the boundary from the finest grid and the simulation starts.

Shows. The investment region under mean reversion, and where the simulation checks it.




Figure 4: The value $F(P, K)$ of the option to invest under a mean-reverting price on the grid with 128 cells in $\ln P$, on logarithmic scales of the price P and the cost K , with the computed boundary $K^*(P)$ (invest when $K \leq K^*(P)$), the long-run price $\bar{p} = 1$ dotted, and the five simulation starts labelled with the PDE value and the simulated mean and standard error.

The boundary rises with the price but flattens, $\ln K^*$ going from 1.548 at $\ln P = -1$ to 2.951 at $\ln P = 1$, because a high price is expected to fall back. At each star the two values agree within one standard error.

The boundary checked by simulating the rule it defines. *Method.* First, Π of Lemma 1 by quadrature against Monte Carlo on exact Ornstein–Uhlenbeck paths (steps of 0.05 years over 250 years, 10^5 paths). Second, the rule "invest when (P, K) enters the computed region", from the finest grid, simulated with exact correlated transitions in steps of 0.01 years over 200 years, 10^5 paths per start, its mean discounted payoff against the PDE value. A stopping rule gives a lower bound [33], [34]: a right boundary puts the mean neither far below the PDE value nor significantly above.

Parameters. Definition 2; starts $(\ln P, \ln K) = (-1, 2), (-0.5, 2.5), (0, 2.75), (0.5, 3.25), (1, 3.25)$.

Result. $\Pi = 19.09, 20.12, 21.44, 23.16, 25.45$ at $\ln P = -1, \dots, 1$, increasing, with simulation z-scores $+1.22, +0.46, -1.44, -0.63, -0.41$. The rule's values are $11.851 \pm 0.003, 8.531 \pm 0.005, 6.689 \pm 0.006, 3.181 \pm 0.006$ and 3.254 ± 0.006 , against PDE values 11.849, 8.528, 6.685, 3.179 and 3.249: z-scores $+0.93, +0.53, +0.73, +0.33, +0.94$.

Error. All five simulated means lie above the PDE values, each within one standard error; the combined z-score, 1.55, is not statistically significant, so the boundary loses no value the simulation can detect. With steps of 0.04 years the estimates drop, one to 2.4 standard errors below, as coarser monitoring invests late. This verifies Lemma 1 and checks the region.

The solvers without a closed form to lean on. *Method.* Algorithms 1, 2 and 3 on



problem (b), from the payoff, to the common test, timed as before.

Parameters. Definition 2; $N = 32$ to 256 cells in P , 4,389 to 271,649 unknowns.

Result. Policies 8, 13, 26, 50; Newton steps 10, 15, 27, 52; sweeps 1886, 597, 1658, 4883. Times: policy iteration 0.022, 0.17, 2.43, 35.1 s; projected SOR 0.14, 0.18, 2.02, 23.7 s; penalty 0.040, 0.38, 6.04, 105.5 s (one run at $N = 256$). Policy iteration is fastest on the coarse grids, projected SOR on the fine ones.

Error. Projected SOR agrees with policy iteration to between 2.7×10^{-10} and 8.8×10^{-12} , the penalty method to 3.2×10^{-10} , against a test of 3.2×10^{-9} . By refinement the largest change at the five starts is 3.6×10^{-3} , 9.1×10^{-4} and 2.3×10^{-4} , orders 1.98 and 2.00. The penalty method again made one or two more solves than policy iteration (Proposition 5); projected SOR's sweeps vary irregularly with the grid.

9. An Option That Expires

Script. Plots the finite-life study against T .

Shows. Proposition 2: values and boundary rise with the life toward the perpetual option.




Figure 5: Left: the investment boundary $X^*(T)$ of the option with life T , by policy iteration with Rannacher stepping on the 128×128 grid with $\Delta t = 1/16$, for T from 1 to 50 years, with the perpetual boundary on the same grid and in closed form; right: the gap to the perpetual solution on the same grid, in value at four points and in the boundary.

The boundary rises from 1.997 at one year to 2.836 at fifty, toward 2.8382 on the same grid. The gaps fall slowly for ten years, then fast; at $X = e^{0.8}$ the gap is flat at first, that point being in the investment region for the two shortest lives.

The finite life approaching the perpetual option. *Method.* Problem (c) for $T = 1, 2, 5, 10, 20, 50$ by Rannacher stepping with policy iteration per step, against the perpetual problem on the same grid, so that the gap measures the life alone.



Parameters. Example 1; $N = 128$; $\Delta t = 0.0625$.

Result. $X^*(T) = 1.997, 2.154, 2.442, 2.650, 2.782$ and 2.836 , against 2.8382 perpetual on this grid. At $X = e^{-1}$ the value grows from 3.0×10^{-6} to 0.0773 , against 0.0792 ; at $X = e^{0.8}$ it is the payoff, 1.2255 , for $T = 1$ and 2 , and 1.2637 at $T = 50$, against 1.2641 .

Error. Values and boundary increase with T at every point, as Proposition 2 states; at $T = 50$ the largest value gap is 1.9×10^{-3} and the boundary gap 2.5×10^{-3} , with each step solved to 2.8×10^{-14} . The grid's own boundary error, -6.2×10^{-3} , is common to both sides. This verifies Proposition 2 on the grid.

Script. Plots the time-stepping study as accuracy against work under Rannacher stepping.

Shows. What each way of solving a time step costs for a given accuracy at expiry.




Figure 6: Sup-norm error at $T = 1$ against the wall time of the whole run, median of three on one thread, for projected SOR, the penalty method, policy iteration and Ikonen–Toivanen splitting under Rannacher stepping on the 128×128 grid, each line running from $\Delta t = 1/4$ to $\Delta t = 1/128$.

All four methods reach nearly the same accuracy at each step size, so the curves differ mainly in cost. Projected SOR and splitting reach an error of 8.9×10^{-6} in under 0.2 s; the penalty method and policy iteration, which factorise a matrix at almost every step, need 3.25 s and 2.43 s for the same accuracy. The orders in the step are those of the time stepping, not of the step solver: implicit Euler 1.00 , Crank–Nicolson 1.99 , Rannacher 2.00 .

Time stepping with four methods. *Method.* Problem (c), $T = 1$, by implicit Euler, Crank–Nicolson and Rannacher stepping (four implicit half steps, then Crank–Nicolson [30]); each step solved by Algorithm 1, 2 or 3 to the common test, or taken by Algorithm 4. Sup-norm errors at T against 2048 Rannacher steps.

Parameters. Example 1; $N = 128$; $\Delta t = 1/4$ to $1/128$.

Result. Observed orders: 0.997 (implicit Euler), 1.993 (Crank–Nicolson), 2.002 (Rannacher); errors at $\Delta t = 1/128$: 1.45×10^{-3} , 2.57×10^{-6} and 8.82×10^{-6} . The exact step



solvers agree, the largest difference being 2.51×10^{-6} against 2.57×10^{-6} ; splitting matches them for implicit Euler and Rannacher and gives 3.26×10^{-6} for Crank–Nicolson. Times: Table 3.

Error. Crank–Nicolson kept second order without smoothing here: the payoff’s kink at $X = 1$ lies in the continuation region, the boundary at expiry being 1.75, and the largest eigenvalue times the step stays moderate on this grid. Rannacher’s start-up, which guards against the loss of order described in [31], cost an error constant about 3.5 times larger. Splitting’s step residual is 1.8×10^{-2} at $\Delta t = 1/4$ and 1.3×10^{-5} at $1/128$, its split complementarity exact. Under implicit Euler the penalty method and policy iteration made the same numbers of solves at every step size, 9 to 136, consistent with Proposition 5.

Source. The time-stepping study of problem (c), $N = 128$, $T = 1$; whole run, median of three, one thread; errors against the 2048-step reference.

Columns. Method; iterations per step from $\Delta t = 1/4$ to $1/128$ under Rannacher stepping; factorisations, time for implicit Euler, Crank–Nicolson and Rannacher, and Rannacher’s error, all at $\Delta t = 1/128$.

Table 3: Work per time step, time of the whole run and error at T for the four ways of taking a step, $N = 128$, $T = 1$, $\Delta t = 1/128$.

method	iterations per step	factorisations	time (s): IE, CN, Rannacher	error
projected SOR	14 to 4 sweeps	0	0.204, 0.169, 0.177	8.88×10^{-6}
penalty method	2.0 to 1.06	136 to 138	3.06, 3.08, 3.25	8.82×10^{-6}
policy iteration	1.83 to 1.05	134 to 136	2.41, 2.41, 2.43	8.82×10^{-6}
splitting	1 solve	1, 1, 2	0.133, 0.143, 0.171	8.79×10^{-6}

The Newton methods need just over one factorisation per step, which is what makes them slow.

10. Choosing a Method

Remark 3 (Which method, when). For a sparse direct solver, one thread and up to about a million unknowns, the measurements support these choices.

1. *Perpetual problems on large grids: projected SOR with Young’s factor.* Fastest on problem (a) at every grid, 2.55 s against 25.7 s and 27.1 s at $N = 512$, and on problem (b) from $N = 128$, 23.7 s against 35.1 s and 105.5 s at $N = 256$; Young’s rule costs about 35 per cent more sweeps than the best factor but avoids the factors that failed.
2. *Coarse grids, or an exact discrete solution: policy iteration.* Fastest on problem (b)



at $N = 32$ and 64 ; its count grows like N , so its cost grows faster than projected SOR's.

3. *The penalty method only with the penalty rule.* It solves policy iteration's systems plus one or two, so it was never faster.
4. *An option that expires: splitting, or projected SOR per step,* 0.13 to 0.20 s against 2.4 s for policy iteration at $\Delta t = 1/128$. Plain Crank–Nicolson sufficed because the kink lies in the continuation region; otherwise Rannacher's start-up is the safe choice, at about 3.5 times the error constant.
5. *Stencil: match the grid to the volatilities,* $h_y/h_x = \sigma_K/\sigma_V$, so that S7 is monotone at every correlation. Elsewhere W restores the guarantee at a cost in accuracy; S9 was the most accurate, without a guarantee, and slowed policy iteration to over 100 policies.

Remark 4 (What this comparison does not cover). The rankings rest on a sparse direct solver, one thread and three problems with constant coefficients in log coordinates.

With an iterative or multigrid inner solver a penalty or policy step costs a few cycles instead of a factorisation, and projected multigrid replaces projected SOR [25], [24]; either could reverse the ranking at large N . With three or more factors a factorisation fills in far more, which favours methods without one. Policy iteration was started from the payoff; a nested start from a coarser grid's exercise set would shorten the walk of the boundary its count measures, and was not tried. The convergence of monotone schemes [29] does not cover the nine-point stencil, whose accuracy is measured here, not proved. The free boundary itself, its regularity and its behaviour near expiry [10], enters only through the closed form and the limit δ_K/δ_V .

11. Conclusion

Investment timing with two random factors is, after a monotone discretisation, a linear complementarity problem with one solution, and all four methods reach it: on the McDonald–Siegel problem their solutions agree to 2.1×10^{-9} , the boundary converges to the closed form at second order, to 2.4×10^{-4} at 512 cells a side, and the reduction confirms it to 6.5×10^{-7} .

What separates the methods is cost and robustness. With a direct solver and one thread, projected SOR was the fastest on the perpetual problems at the larger grids, because the discount keeps its iteration contracting and a sweep is cheap. Policy iteration from the payoff moves the boundary about one cell per iteration, so its finite count grows with the grid. The penalty method is policy iteration with a finite penalty, and needs that penalty set by a rule. Monotonicity guarantees convergence, not accuracy. For an expiring option, splitting was as accurate as the exact step solvers in 0.13 to 0.17 s against 2.4 s. Under mean reversion the simulated rule recovered the PDE value within one standard error at five points. Remark 3 holds under these conditions.



References

- [1] McDonald, R., & Siegel, D. (1986). The value of waiting to invest. *The Quarterly Journal of Economics*. <https://doi.org/10.2307/1884175>
- [2] Dixit, A. K., & Pindyck, R. S. (1994). *Investment under Uncertainty*. Princeton University Press. <https://doi.org/10.1515/9781400830176>
- [3] Dixit, A. (1989). Entry and exit decisions under uncertainty. *Journal of Political Economy*. <https://doi.org/10.1086/261619>
- [4] Brennan, M. J., & Schwartz, E. S. (1985). Evaluating natural resource investments. *The Journal of Business*. <https://doi.org/10.1086/296288>
- [5] Schwartz, E. S. (1997). The stochastic behavior of commodity prices: implications for valuation and hedging. *The Journal of Finance*. <https://doi.org/10.1111/j.1540-6261.1997.tb02721.x>
- [6] Bensoussan, A., & Lions, J.-L. (1982). *Applications of Variational Inequalities in Stochastic Control*. North-Holland. [https://doi.org/10.1016/s0168-2024\(08\)x7013-4](https://doi.org/10.1016/s0168-2024(08)x7013-4)
- [7] Jaillet, P., Lamberton, D., & Lapeyre, B. (1990). Variational inequalities and the pricing of American options. *Acta Applicandae Mathematicae*. <https://doi.org/10.1007/bf00047211>
- [8] Peskir, G., & Shiryaev, A. (2006). *Optimal Stopping and Free-Boundary Problems*. Birkhäuser. <https://doi.org/10.1007/978-3-7643-7390-0>
- [9] Pham, H. (2009). *Continuous-time Stochastic Control and Optimization with Financial Applications*. Springer. <https://doi.org/10.1007/978-3-540-89500-8>
- [10] Dewynne, J. N., Howison, S. D., Rupf, I., & Wilmott, P. (1993). Some mathematical results in the pricing of American options. *European Journal of Applied Mathematics*. <https://doi.org/10.1017/s0956792500001194>
- [11] Brennan, M. J., & Schwartz, E. S. (1977). The valuation of American put options. *The Journal of Finance*. <https://doi.org/10.2307/2326779>
- [12] Cottle, R. W., Pang, J.-S., & Stone, R. E. (2009). *The Linear Complementarity Problem*. SIAM. <https://doi.org/10.1137/1.9780898719000>
- [13] Cryer, C. W. (1971). The solution of a quadratic programming problem using systematic overrelaxation. *SIAM Journal on Control*. <https://doi.org/10.1137/0309028>
- [14] Zvan, R., Forsyth, P. A., & Vetzal, K. R. (1998). Penalty methods for American options with stochastic volatility. *Journal of Computational and Applied Mathematics*. [https://doi.org/10.1016/s0377-0427\(98\)00037-5](https://doi.org/10.1016/s0377-0427(98)00037-5)
- [15] Forsyth, P. A., & Vetzal, K. R. (2002). Quadratic convergence for valuing American options using a penalty method. *SIAM Journal on Scientific Computing*. <https://doi.org/10.1137/s1064827500382324>



- [16] Witte, J. H., & Reisinger, C. (2011). A penalty method for the numerical solution of Hamilton-Jacobi-Bellman (HJB) equations in finance. *SIAM Journal on Numerical Analysis*. <https://doi.org/10.1137/100797606>
- [17] Howard, R. A. (1960). *Dynamic Programming and Markov Processes*. Technology Press of MIT and Wiley.
- [18] Bokanowski, O., Maroso, S., & Zidani, H. (2009). Some convergence results for Howard's algorithm. *SIAM Journal on Numerical Analysis*. <https://doi.org/10.1137/08073041x>
- [19] Hintermüller, M., Ito, K., & Kunisch, K. (2002). The primal-dual active set strategy as a semismooth Newton method. *SIAM Journal on Optimization*. <https://doi.org/10.1137/s1052623401383558>
- [20] Reisinger, C., & Witte, J. H. (2012). On the use of policy iteration as an easy way of pricing American options. *SIAM Journal on Financial Mathematics*. <https://doi.org/10.1137/110823328>
- [21] Huang, Y., Forsyth, P. A., & Labahn, G. (2012). Combined fixed point and policy iteration for Hamilton-Jacobi-Bellman equations in finance. *SIAM Journal on Numerical Analysis*. <https://doi.org/10.1137/100812641>
- [22] Ikonen, S., & Toivanen, J. (2004). Operator splitting methods for American option pricing. *Applied Mathematics Letters*. <https://doi.org/10.1016/j.aml.2004.06.010>
- [23] Ikonen, S., & Toivanen, J. (2009). Operator splitting methods for pricing American options under stochastic volatility. *Numerische Mathematik*. <https://doi.org/10.1007/s00211-009-0227-5>
- [24] Toivanen, J., & Oosterlee, C. W. (2012). A projected algebraic multigrid method for linear complementarity problems. *Numerical Mathematics: Theory, Methods and Applications*. <https://doi.org/10.4208/nmtma.2011.m12si05>
- [25] Clarke, N., & Parrott, K. (1999). Multigrid for American option pricing with stochastic volatility. *Applied Mathematical Finance*. <https://doi.org/10.1080/135048699334528>
- [26] Barles, G., & Souganidis, P. E. (1991). Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Analysis*. <https://doi.org/10.3233/asy-1991-4305>
- [27] Bonnans, J. F., & Zidani, H. (2003). Consistency of generalized finite difference schemes for the stochastic HJB equation. *SIAM Journal on Numerical Analysis*. <https://doi.org/10.1137/s0036142901387336>
- [28] Ma, K., & Forsyth, P. A. (2016). An unconditionally monotone numerical scheme for the two-factor uncertain volatility model. *IMA Journal of Numerical Analysis*. <https://doi.org/10.1093/imanum/drw025>



- [29] Kushner, H. J., & Dupuis, P. (2001). *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer.
<https://doi.org/10.1007/978-1-4613-0007-6>
- [30] Rannacher, R. (1984). Finite element solution of diffusion problems with irregular data. *Numerische Mathematik*. <https://doi.org/10.1007/bf01390130>
- [31] Pooley, D. M., Vetzal, K. R., & Forsyth, P. A. (2003). Convergence remedies for non-smooth payoffs in option pricing. *The Journal of Computational Finance*.
<https://doi.org/10.21314/jcf.2003.101>
- [32] Achdou, Y., & Pironneau, O. (2005). *Computational Methods for Option Pricing*. SIAM. <https://doi.org/10.1137/1.9780898717495>
- [33] Longstaff, F. A., & Schwartz, E. S. (2001). Valuing American options by simulation: a simple least-squares approach. *Review of Financial Studies*.
<https://doi.org/10.1093/rfs/14.1.113>
- [34] Haugh, M. B., & Kogan, L. (2004). Pricing American options: a duality approach. *Operations Research*. <https://doi.org/10.1287/opre.1030.0070>



Zamrik Quant Research

The research arm of Zamrik Quant Lab

About the Desk

Optimal Control & Decision Desk

Zamrik Quant Research


What to do under uncertainty: optimal control and HJB, mean-field games, agent-based models, irreversible decisions, betting, lotteries and insurance.


Most of this desk's questions have the same shape: someone must act now, without knowing what comes next, and live with the result. The Optimal Control & Decision desk turns such questions into Hamilton–Jacobi–Bellman equations, optimal stopping and singular control problems, and solves them.

When several parties optimise at once, the problem becomes a game: Dynkin games and wars of attrition between two or three players, and mean field games, in which a large population and its representative member shape each other through coupled Hamilton–Jacobi–Bellman and Fokker–Planck equations. Where the equations give out, agent-based models take over, and the desk says plainly which tool fits which case.


The applications are concrete: when to build an irreversible plant, how a company should pay dividends, how to run a home battery against a time-of-use tariff, when a wind turbine should turn into the wind, how a central bank should answer stagflation, and how much to bet when an edge is small. Its notebooks and essays carry the same reasoning to retirement, market timing, lotteries, insurance and prize competitions, each worked through to a decision a reader could actually make.

The desk is supervised by **Dr. T. Zamrik**, Head of Zamrik Quant Research, PhD, Imperial College London (2011).





Confocal conics: $\frac{x^2}{a^2-\lambda} + \frac{y^2}{b^2-\lambda} = 1$, ellipses for $\lambda < b^2$, hyperbolae for $b^2 < \lambda < a^2$, foci $(\pm\sqrt{a^2-b^2}, 0)$



ZAMRIK QUANT RESEARCH

The research arm of Zamrik Quant Lab



TR-2026-08169923
zamrik.com