


WORKING PAPER

Fredholm Methods for Optimal Stopping under Lévy Noise

A Wiener-Hopf Cascade for Real Options under Lévy Noise

Supervised by Dr. T. Zamrik

21-MAY-2025



Sample paths of stochastic logistic growth: $dX_t = r X_t(1 - X_t/K) dt + \sigma X_t dW_t$, $r = 9.48$, $K = 1$, $\sigma = 0.27$, $X_0 = 0.02$



WP-2026-62242748
zamrik.com



Zamrik Quant Research

The research arm of Zamrik Quant Lab

Copyright

© 2025 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 (2025). ‘Fredholm Methods for Optimal Stopping under Lévy Noise: A Wiener-Hopf Cascade for Real Options under Lévy Noise’. *Zamrik Quant Research Working Paper WP-2026-62242748*. Available at <https://zamrik.com/research-items/wp-2026-62242748/>.

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 paper 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.



WORKING PAPER

Fredholm Methods for Optimal Stopping under Lévy Noise

A Wiener-Hopf Cascade for Real Options under Lévy Noise

Supervised by Dr. T. Zamrik¹

21-MAY-2025

WP-2026-62242748

Abstract. We develop a Fredholm-operator framework for optimal stopping of Lévy processes with exponential tails, built on the Wiener–Hopf factorisation of the resolvent. The resolvent symbol factorises as $q - \psi(\xi) = \phi_q^+(\xi) \phi_q^-(\xi)$, where $\phi_q^\pm$ are the reciprocal transforms of the supremum and the infimum of the process at an independent exponential time. For payoffs that increase with the state, the variational inequality is solved by a two-level cascade. A Fredholm convolution equation of the first kind, whose kernel is the law of the supremum, gives an averaging function. The stopping region is the half-line above the root of that function, and the value is the supremum-average of its positive part. Under Kou double-exponential jumps both levels are explicit, and the investment threshold for the payoff $(P - I)^+$ is $P^* = I \mathbb{E}[e^S] = I \alpha_1 \alpha_2 (\eta_+ - 1) / (\eta_+ (\alpha_1 - 1)(\alpha_2 - 1))$, where α_1, α_2 are the two positive roots of $\kappa(r) = q$. We fit the Kou model by exact maximum likelihood to 1,567 EU ETS primary-auction clearing prices from 2018-01-08 to 2025-05-20 (EEX auction results as published by the Instrat Foundation). A probability-integral-transform test does not reject the Kou fit (Kolmogorov–Smirnov $p = 0.98$) and rejects a Gaussian fit to the same prices ($p = 0.005$). With the price expected to grow at the euro risk-free rate of 2.73 %, the threshold multiple P^*/I falls from 4.53 at a 7 % hurdle rate to 2.29 at 18 %. The McDonald–Siegel threshold of a geometric Brownian motion fitted to the same prices lies 0.2–0.3 % above the Kou threshold at every hurdle rate. On these data the jump-aware threshold is therefore not stable across hurdle rates: the jumps move it by less than half a per cent, and the hurdle rate moves it by a factor of 1.98. Each threshold is checked by a grid solution of the variational inequality and by Monte Carlo simulation of the stopping rule.

WP-2026-62242748 · Optimal Control & Decision Desk · Working Paper · Published
21-MAY-2025 · Reading time: about 38 min

1. Introduction

A firm that must commit to irreversible capital, such as green-energy generation, transmission or storage, chooses *when* to install, not *whether*. The classical real-options framework ([1], [2]) answers this question when the underlying price follows a geometric Brownian motion: invest the first time the price crosses a threshold set by the drift, the volatility and the discount rate. The carbon price that drives the value of green capacity in Europe is the price of an EU Emissions Trading System allowance (EUA), and the EUA auction record shows that a geometric Brownian motion does not describe it well. Its daily moves

¹Correspondence: Optimal Control & Decision Desk, Zamrik Quant Research, info@zamrik.com.



have fat tails, and the largest of them are discrete jumps (Section 7). This paper asks two questions. Does the investment threshold change when jumps are modelled? And how much of the threshold is a property of the data rather than of the firm's hurdle rate?

1.1 Motivation: irreversibility meets jump risk

Between 2018-01-08 and 2025-05-20 the EUA primary-auction clearing price ranged from 7.68 €/tCO₂ (2018-01-09) to 97.51 €/tCO₂ (2022-02-08) ([3]). The one-business-day log-returns between auctions have excess kurtosis 2.83. Two auction-to-auction moves exceed 15 % in absolute value, and both are falls: -25.3 % on 2020-03-19 and -22.5 % on 2022-03-03. A real-options analysis that fits a single volatility to such a series may misplace the trigger, because the value of waiting depends on the whole law of the price's running maximum, not only on its variance. Whether it does so here is a quantitative question, and Section 7 answers it.

1.2 The mathematical problem

The optimal-stopping problem under Lévy noise is governed by a non-local variational inequality on the real line,

$$\max(\mathcal{L}V - qV, g - V) = 0, \quad (1.1)$$

where $\mathcal{L}$ is the integro-differential generator of the Lévy process, q the discount rate, and g the payoff. Finite-difference schemes for the partial integro-differential equation discretise both the differential and the non-local part on a fixed grid; a dense quadrature of the jump integral costs $O(N^2)$ work per step ([4]; [5] for the variance gamma process). Such schemes produce a number for a given parameter set, but no closed-form expression for the value function or the threshold.

The Wiener–Hopf factorisation of a Lévy process gives such expressions ([6], [7]). This paper organises it into a constructive scheme for one-sided problems. The resolvent factorises into the resolvents of the running supremum and the running infimum ([8], Theorem 1). For a payoff that increases with the state, the value is obtained in two steps. First, solve the *Fredholm integral equation of the first kind*

$$\int_{s \geq 0} Q(x + s) \mathbb{P}(S \in ds) = G(x), \quad x \in \mathbb{R}, \quad (1.2)$$

whose kernel is the law of the supremum S of the process at an independent exponential time. Then average the positive part of its solution Q against the same law. The stopping boundary is the root of Q . We call the two steps the *cascade*.

1.3 Contribution and roadmap

The paper establishes five claims.

1. **Wiener–Hopf factorisation (Thm 1).** The resolvent symbol of a Lévy process



with exponential tails factors as $q - \psi = \phi_q^+ \phi_q^-$, with $\phi_q^\pm$ the reciprocal transforms of the supremum and the infimum at an exponential time, holomorphic and zero-free in the upper and lower half-planes. This is the published factorisation of [8] in Fourier form.

2. **The two-level cascade (Thms 2–3, Cor 2).** For a non-decreasing payoff G^+ , the value is $V = \mathbb{E}[Q(\cdot + S)^+]$, where the averaging function Q solves the first-kind Fredholm equation above. The stopping region is the half-line above the root of Q , and continuous fit always holds. For payoffs of exponential type, and for any smooth payoff when the factor is rational (Kou), Q is explicit.
3. **The Kou threshold (Thm 4).** For the green-capacity timing problem under a Kou jump-diffusion, $P^* = I \Phi_q^+(1)$, set by the *two* positive roots of $\kappa(r) = q$. The threshold falls strictly as the hurdle rate rises, tends to infinity as the hurdle rate falls to the expected growth rate of the price, and tends to the McDonald–Siegel threshold as the jump intensity tends to zero.
4. **EU ETS fit (§7.1–7.2).** An exact-likelihood Kou fit to 1,566 auction-to-auction returns. Jumps carry 39 % of the return variance, and the fit passes a probability-integral-transform test that rejects the Gaussian model.
5. **Threshold against McDonald–Siegel (§7.3–7.4).** Across hurdle rates from 7 % to 18 %, the Kou threshold multiple falls from 4.53 to 2.29. The McDonald–Siegel multiple is 0.2–0.3 % higher at every rate. Both are checked by a grid solve of the variational inequality and by Monte Carlo.

§2 fixes notation. §3 recasts the variational inequality as a resolvent fixed-point equation. §4 states the factorisation and the cascade. §5 describes the computations and the two independent checks. §6 specialises to Kou jumps and derives the threshold. §7 fits the model to EU ETS data and compares the threshold with McDonald–Siegel. §8 concludes. Appendix A gives the comparison-principle proof, Appendix B the explicit Kou factors, and Appendix C the algorithms.

2. Preliminaries

Throughout the paper, $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ is a filtered probability space supporting a Lévy process X_t , and $\mathbf{e}(q)$ denotes an exponential random time of rate $q > 0$ independent of X . The discount rate $q > 0$ and the payoff $g : \mathbb{R} \rightarrow \mathbb{R}_+$ are fixed.

2.1 Lévy processes with exponential tails

Definition 1 (Lévy process). X_t is a càdlàg $(\mathcal{F}_t)$ -adapted process with stationary independent increments and $X_0 = 0$, with Lévy triplet (b, σ^2, ν) , $\sigma \geq 0$. Its characteristic function is $\mathbb{E}[e^{i\xi X_t}] = e^{t\psi(\xi)}$, where the *characteristic exponent* has the Lévy–Khintchine form

$$\psi(\xi) = ib\xi - \frac{1}{2}\sigma^2\xi^2 + \int_{\mathbb{R}} (e^{i\xi y} - 1 - i\xi y \mathbb{1}_{|y| \leq 1}) \nu(dy),$$



with ν the Lévy measure on $\mathbb{R} \setminus \{0\}$. Where it is finite, the *cumulant* is $\kappa(r) := \log \mathbb{E}[e^{rX_1}] = \psi(-ir)$.

We adopt two standing assumptions.

Assumption 1 (A1, two-sided exponential tails). There are $\eta_+ > 1$ and $\eta_- > 0$ such that $\int_{y>1} e^{ry} \nu(dy) < \infty$ for every $r < \eta_+$ and $\int_{y<-1} e^{ry} \nu(dy) < \infty$ for every $r > -\eta_-$. Then κ is finite and real-analytic on $(-\eta_-, \eta_+)$, and ψ continues holomorphically to the strip $\{\xi \in \mathbb{C} : -\eta_+ < \Im \xi < \eta_-\}$.

Assumption 2 (A2, discounting). $q > 0$. For payoffs that grow like e^x (the call of §6), in addition $q > \kappa(1)$, that is $\mathbb{E}[e^{X_1}] < e^q$.

Since $|\mathbb{E}[e^{i\xi X_1}]| \leq 1$, we have $\Re \psi(\xi) \leq 0$ on $\mathbb{R}$, so $\Re(q - \psi(\xi)) \geq q > 0$ there: the resolvent symbol of §3 never vanishes on the real line. The Kou model of §6 satisfies A1 with its two jump rates.

2.2 The generator and the growth class

Definition 2 (Generator). For $\varphi \in C^2$ with $\varphi, \varphi', \varphi''$ of at most exponential growth inside the strip of A1,

$$\mathcal{L}\varphi(x) = b\varphi'(x) + \frac{1}{2}\sigma^2\varphi''(x) + \int_{\mathbb{R}} [\varphi(x+y) - \varphi(x) - y\varphi'(x)\mathbb{1}_{|y|\leq 1}] \nu(dy).$$

On exponentials, $\mathcal{L}e^{i\xi x} = \psi(\xi)e^{i\xi x}$ for $\xi \in \mathbb{R}$, and $\mathcal{L}e^{rx} = \kappa(r)e^{rx}$ for $r \in (-\eta_-, \eta_+)$.

We write $\mathcal{G}$ for the class of continuous functions with $|\varphi(x)| \leq C(1+e^x)$: bounded at $-\infty$, at most exponential of order one at $+\infty$.

Lemma 1 (Generator bound). Under A1, for $\varphi \in C_b^2(\mathbb{R})$,

$$\|\mathcal{L}\varphi\|_{\infty} \leq |b|\|\varphi'\|_{\infty} + \frac{1}{2}\left(\sigma^2 + \int_{|y|\leq 1} y^2 \nu(dy)\right)\|\varphi''\|_{\infty} + 2\nu(\{|y| > 1\})\|\varphi\|_{\infty}.$$

Proof. For $|y| \leq 1$, Taylor's theorem with integral remainder gives $|\varphi(x+y) - \varphi(x) - y\varphi'(x)| \leq \frac{1}{2}y^2\|\varphi''\|_{\infty}$. For $|y| > 1$ the integrand is bounded by $2\|\varphi\|_{\infty}$, and $\nu(\{|y| > 1\}) < \infty$ for every Lévy measure. Adding the drift and diffusion terms gives the bound. $\square$

2.3 Optimal stopping and the variational inequality

For $x \in \mathbb{R}$, define the value of the optimal-stopping problem

$$V(x) = \sup_{\tau \in \mathcal{T}} \mathbb{E}_x[e^{-q\tau}g(X_{\tau})], \quad (2.1)$$

where $\mathcal{T}$ is the set of $(\mathcal{F}_t)$ -stopping times, $e^{-q\tau}g(X_{\tau}) := 0$ on $\{\tau = \infty\}$, and $\mathbb{E}_x$ denotes expectation under $X_0 = x$. Standard arguments ([9], [8] Theorems 5–6) show that V is the unique viscosity solution in $\mathcal{G}$ of the following non-local variational inequality.



Definition 3 (Variational inequality of optimal stopping). A function $V \in \mathcal{G}$ satisfies the variational inequality if

$$\max(\mathcal{L}V(x) - qV(x), g(x) - V(x)) = 0 \quad \text{for every } x \in \mathbb{R},$$

in the viscosity sense.

Two regions partition the real line: the *continuation region*

$$\mathcal{C} = \{x \in \mathbb{R} : V(x) > g(x)\}, \tag{2.2}$$

on which $\mathcal{L}V - qV = 0$, and the *stopping region* $\mathcal{D} = \mathcal{C}^c$, on which $V = g$. Continuous fit, $V|_{\partial\mathcal{C}} = g|_{\partial\mathcal{C}}$, always holds. Smooth fit holds exactly when 0 is regular for the half-line on the stopping side, in particular whenever $\sigma > 0$; otherwise the value may have a kink at the boundary ([10]; [8], Theorem 7(ii)).

2.4 Payoff class

The payoffs of §§6–7 are the perpetual call $g(x) = (e^x - K)^+$ (green-capacity timing) and the perpetual put $g(x) = (K - e^x)^+$. Both lie in $\mathcal{G}$ and are Lipschitz in the price $P = e^x$; the call needs $q > \kappa(1)$ (A2) for a finite value. §4 treats the larger class $g = G^+$ with G non-decreasing, the *one-sided* problems, whose stopping region is a half-line above a threshold.

3. Variational Inequality and Resolvent Reformulation

The variational inequality (Def. 3) is a free-boundary problem with two unknowns: the value function V and the partition $\{\mathcal{C}, \mathcal{D}\}$. This section writes it as a fixed-point equation for the resolvent of $\mathcal{L}$ and records the comparison principle (Lem. 3) that secures uniqueness; the proof is in Appendix A.

3.1 Free-boundary structure

Within the continuation region $\mathcal{C}$ the value function satisfies $(q - \mathcal{L})V = 0$; within the stopping region $\mathcal{D}$ it satisfies $V = g$. Because $\mathcal{L}$ is non-local, the equation on $\mathcal{C}$ involves the values of V on $\mathcal{D}$ wherever jumps reach across the boundary. A solution must therefore match the payoff on the far side of the boundary, not just at the boundary point.

The usual approach is to guess the structure of $\mathcal{C}$, say a half-line $(-\infty, b)$ for a call-type payoff, solve $(q - \mathcal{L})V = 0$ on it, and fix b by the fit conditions. For one-sided payoffs, the cascade of §4 derives the half-line structure instead of assuming it.



3.2 Resolvent reformulation

Definition 4 (Resolvent operator). For $q > 0$,

$$\mathcal{R}_q f(x) = \mathbb{E}_x \left[\int_0^\infty e^{-qt} f(X_t) dt \right] = \frac{1}{q} \mathbb{E}[f(x + X_{\mathbf{e}(q)})],$$

a bounded operator on $L^\infty(\mathbb{R})$ with norm $1/q$. It acts on exponentials by $\mathcal{R}_q e^{i\xi x} = e^{i\xi x}/(q - \psi(\xi))$ and $\mathcal{R}_q e^{rx} = e^{rx}/(q - \kappa(r))$ when $\kappa(r) < q$; under A2 it maps $\mathcal{G}$ into itself ([8], Definition 1 and Lemma 2).

Lemma 2 (Resolvent fixed-point equation). Assume A1–A2. Let the value function $V \in \mathcal{G}$ be C^1 on $\mathbb{R}$ and C^2 off $\partial\mathcal{C}$ with V'' locally bounded. Then $h := (q - \mathcal{L})V$ vanishes on $\mathcal{C}$, is non-negative on $\mathcal{D}$, and

$$V = \mathcal{R}_q[\mathbb{1}_{\mathcal{D}}(qg - \mathcal{L}V)], \quad \mathcal{D} = \{V = g\},$$

a fixed-point equation for V whose source is supported on the stopping region.

Proof. On $\mathcal{C}$, $(q - \mathcal{L})V = 0$. On the interior of $\mathcal{D}$, $V = g$, so $h = qg - \mathcal{L}V$, and $h \geq 0$ by the variational inequality. Since V'' is locally bounded and ν integrates e^y on $\{y > 1\}$ (A1), Lemma 1, applied to V multiplied by a smooth cut-off, shows that h is a locally bounded function. Under the regularity assumed, the Itô–Meyer formula applies to $e^{-qt}V(X_t)$, and $e^{-qt}V(X_t) + \int_0^t e^{-qs}h(X_s)ds$ is a local martingale. Its localising sequence can be removed because $V \in \mathcal{G}$ and $q > \kappa(1)$. As $t \rightarrow \infty$, $\mathbb{E}_x[e^{-qt}V(X_t)] \leq C(e^{-qt} + e^x e^{-(q-\kappa(1))t}) \rightarrow 0$. Hence $V(x) = \mathbb{E}_x \int_0^\infty e^{-qs}h(X_s)ds = \mathcal{R}_q h(x)$, and $h = \mathbb{1}_{\mathcal{D}}h$ up to the null set $\partial\mathcal{C}$, which the process spends zero time in. $\square$

The source $\mathbb{1}_{\mathcal{D}}(qg - \mathcal{L}V)$ is supported on the stopping region, and $\mathcal{R}_q$ acts through the symbol $(q - \psi(\xi))^{-1}$. Theorem 1 of §4 splits this symbol into a factor carried by the supremum and a factor carried by the infimum. For one-sided problems only one of the two is needed.

3.3 Comparison principle and uniqueness

The fixed-point reformulation does not by itself secure uniqueness of V . Uniqueness comes from the following comparison principle, proved in Appendix A by Crandall–Ishii–Lions doubling of variables adapted to the non-local generator.

Lemma 3 (Comparison principle). Assume A1–A2. Let $\underline{V}$ be an upper-semicontinuous viscosity subsolution and $\overline{V}$ a lower-semicontinuous viscosity supersolution of the variational inequality (Def. 3), both in $\mathcal{G}$. Then $\underline{V} \leq \overline{V}$ on $\mathbb{R}$.

Proof. Appendix A (Theorem 5). The argument combines the doubling of variables of Crandall–Ishii–Lions [11] with the non-local maximum principle of Jakobsen–Karlsen [12] and the small-jump/large-jump split of Barles–Imbert [13], as in [8], Theorem 5. $\square$



Corollary 1 (Uniqueness). Under A1–A2, the variational inequality (Def. 3) has at most one viscosity solution in $\mathcal{G}$. The optimal-stopping value function V is a viscosity solution ([9]), so it is the unique one.

Proof. Apply Lemma 3 to $(\underline{V}, \overline{V}) = (V_1, V_2)$ and to (V_2, V_1) for two solutions V_1, V_2 . $\square$

4. The Wiener-Hopf Cascade

This section is the centre of the paper. Theorem 1 factorises the resolvent symbol into a supremum part and an infimum part. Theorem 2 solves every one-sided problem with the supremum part alone, in two levels: an averaging function from a Fredholm equation, then an average of its positive part. Theorem 3 makes the first level explicit for exponential-type payoffs and for rational factors.

4.1 Wiener–Hopf factorisation of the resolvent symbol

Write $S := \sup_{0 \leq t \leq e(q)} X_t \geq 0$ and $\underline{S} := \inf_{0 \leq t \leq e(q)} X_t \leq 0$ for the supremum and the infimum at the independent exponential time. Let $\Phi_q^+(r) := \mathbb{E}[e^{rS}]$ and $\Phi_q^-(r) := \mathbb{E}[e^{r\underline{S}}]$ be their transforms, finite at least for $r \leq 0$ and $r \geq 0$ respectively ([8], Definition 2).

Theorem 1 (Wiener–Hopf factorisation of the resolvent symbol). Assume A1 and $q > 0$. For $\xi \in \mathbb{R}$,

$$q - \psi(\xi) = \phi_q^+(\xi) \phi_q^-(\xi), \quad \phi_q^+(\xi) := \frac{\sqrt{q}}{\mathbb{E}[e^{i\xi S}]}, \quad \phi_q^-(\xi) := \frac{\sqrt{q}}{\mathbb{E}[e^{i\xi \underline{S}}]}.$$

The factor ϕ_q^+ extends to a holomorphic, zero-free function on the upper half-plane $\{\Im \xi > 0\}$, continuous up to the real axis, and ϕ_q^- likewise on the lower half-plane; $\phi_q^\pm(0) = \sqrt{q}$. In operator form, $\mathcal{R}_q = q \mathcal{R}_q^+ \mathcal{R}_q^-$, where $\mathcal{R}_q^+ f(x) = q^{-1} \mathbb{E}[f(x + S)]$ and $\mathcal{R}_q^- f(x) = q^{-1} \mathbb{E}[f(x + \underline{S})]$.

Proof. By the Wiener–Hopf factorisation of the published resolvent paper ([8], Theorem 1), S and $X_{e(q)} - S$ are independent and $X_{e(q)} - S$ has the law of $\underline{S}$. Hence $\mathbb{E}[e^{i\xi X_{e(q)}}] = \mathbb{E}[e^{i\xi S}] \mathbb{E}[e^{i\xi \underline{S}}]$. The left side equals $q \int_0^\infty e^{-qt} e^{t\psi(\xi)} dt = q/(q - \psi(\xi))$, because $\Re \psi \leq 0$. Inverting both sides gives the factorisation, and the operator form is the same identity applied to $f(x + X_{e(q)})$ (Definition 4). Since $S \geq 0$, $\xi \mapsto \mathbb{E}[e^{i\xi S}]$ is holomorphic and bounded on $\Im \xi > 0$ and continuous up to the axis. It has no zeros there, because S is infinitely divisible with Lévy measure on $(0, \infty)$, so the transform is the exponential of $\int_0^\infty \int_{(0, \infty)} (e^{i\xi x} - 1) t^{-1} e^{-qt} \mathbb{P}(X_t \in dx) dt$ ([6], Chapter 6). The same argument applies to $\underline{S}$ on $\Im \xi < 0$, and $\mathbb{E}[e^0] = 1$ gives the normalisation. $\square$

Under A1 the factors continue further: ϕ_q^+ to $\Im \xi > -\alpha_1$ and ϕ_q^- to $\Im \xi < \beta_1$, where $\alpha_1 > 0$ and $-\beta_1 < 0$ are the roots of $\kappa(r) = q$ inside the strip, whenever these exist ([8], Lemma 3). Figure 1 shows the symbol for the fitted Kou model of §7, with the zeros of the two factors on either side of the real axis.




Figure 1: The Kou symbol at the EU ETS fit of Section 7 (Table 1, with the drift set by $\kappa(1) = r$, Section 7.2) and $q = 0.10$: heat map of $\log_{10} |q - \psi(\xi)|$ on the complex ξ -plane. The zeros of $q - \psi$ are at $\xi = -i\alpha_1, -i\alpha_2$ (below the axis: zeros of ϕ_q^+) and at $\xi = i\beta_1, i\beta_2$ (above: zeros of ϕ_q^-); the poles are at $\xi = -i\eta_+$ and $\xi = i\eta_-$. Each factor is zero-free on its own half-plane, as Theorem 1 requires. Teal: the real axis, where the factorisation holds, and (dashed) the edges of the strip of Assumption 1.

4.2 The two-level cascade for one-sided problems

Theorem 2 (Two-level cascade). Assume A1, $q > 0$, and let $G : \mathbb{R} \rightarrow \mathbb{R}$ be continuous and non-decreasing, $g := G^+$. Suppose there is a continuous, strictly increasing $Q : \mathbb{R} \rightarrow \mathbb{R}$, the *averaging function*, with $\mathbb{E}[Q(x + S)] < \infty$ and

$$\mathbb{E}[Q(x + S)] = G(x) \quad \text{for every } x \in \mathbb{R}, \quad \text{that is} \quad q \mathcal{R}_q^+ Q = G,$$

and that Q has a root b^* : $Q < 0$ on $(-\infty, b^*)$ and $Q > 0$ on (b^*, ∞) . Then: (i) the value is $V(x) = \mathbb{E}[Q(x + S)^+] = q \mathcal{R}_q^+[Q^+](x)$, and $\tau^* = \inf\{t \geq 0 : X_t \geq b^*\}$ is optimal; (ii) if $\mathbb{P}(S > y) > 0$ for every $y \geq 0$, the stopping region is exactly $\mathcal{D} = [b^*, \infty)$; (iii) continuous fit $V(b^*) = g(b^*)$ always holds. If Q is C^1 with $Q'(b^*) > 0$, smooth fit $V'(b^*-) = G'(b^*)$ holds if and only if $\mathbb{P}(S = 0) = 0$, that is, if and only if 0 is regular for $(0, \infty)$; in particular whenever $\sigma > 0$.

Proof. This is the argument of Novikov and Shiryaev [14] in its continuous-time form [15], transposed to the supremum; [8], Theorem 7 gives the infimum version. Let $H(x) := \mathbb{E}[Q(x + S)^+]$.

$H \geq g$. By Jensen's inequality for $t \mapsto t^+$, $H(x) \geq (\mathbb{E}[Q(x + S)])^+ = G(x)^+ = g(x)$.

H dominates every stopping rule. Let τ be a stopping time. Since $\mathbf{e}(q)$ is independent of X , $\mathbb{E}_x[e^{-q\tau} H(X_\tau)] = \mathbb{E}_x[H(X_\tau) \mathbb{1}_{\{\tau < \mathbf{e}(q)\}}]$. By the strong Markov property at τ and the lack of memory of $\mathbf{e}(q)$, on $\{\tau < \mathbf{e}(q)\}$ the supremum of X over $[\tau, \mathbf{e}(q)]$ is $X_\tau + S'$, with



S' independent of $\mathcal{F}_\tau$ and distributed as S . Hence, since Q is non-decreasing,

$$\mathbb{E}_x[e^{-q\tau}H(X_\tau)] = \mathbb{E}_x\left[Q(\sup_{\tau \leq t \leq \mathbf{e}(q)} X_t)^+ \mathbb{1}_{\{\tau < \mathbf{e}(q)\}}\right] \leq \mathbb{E}_x\left[Q(\sup_{0 \leq t \leq \mathbf{e}(q)} X_t)^+\right] = H(x).$$

Together with $H \geq g$ this gives $\mathbb{E}_x[e^{-q\tau}g(X_\tau)] \leq H(x)$ for every τ , so $V \leq H$.

Equality at τ^* . Before τ^* the path stays below b^* . On $\{\tau^* \geq \mathbf{e}(q)\}$ the supremum over $[0, \mathbf{e}(q)]$ is at most b^* , where Q^+ vanishes. On $\{\tau^* < \mathbf{e}(q)\}$ the suprema over $[0, \mathbf{e}(q)]$ and over $[\tau^*, \mathbf{e}(q)]$ coincide, because $X_{\tau^*} \geq b^*$. So the inequality above is an equality for $\tau = \tau^*$. For $y \geq b^*$, $Q(y+S) \geq Q(y) \geq 0$, so $H(y) = \mathbb{E}[Q(y+S)] = G(y) \geq 0$, that is $H(y) = g(y)$. Since $X_{\tau^*} \geq b^*$, $H(x) = \mathbb{E}_x[e^{-q\tau^*}g(X_{\tau^*})] \leq V(x)$. Hence $V = H$ and (i) holds.

(ii) We have just shown $V = g$ on $[b^*, \infty)$. For $x < b^*$, $V(x) - G(x) = \mathbb{E}[Q(x+S)^-] > 0$, because $\mathbb{P}(S < b^* - x) > 0$ (the exponential time is small with positive probability). And $V(x) > 0$, because $\mathbb{P}(x+S > b^*) > 0$ by hypothesis and $Q > 0$ above b^* . So $V > \max(G, 0) = g$ on $(-\infty, b^*)$.

(iii) V is continuous by dominated convergence, and $V(b^*) = G(b^*) = g(b^*)$. For $x < b^*$, $V'(x) = \mathbb{E}[Q'(x+S) \mathbb{1}_{\{x+S > b^*\}}]$, which tends to $\mathbb{E}[Q'(b^*+S) \mathbb{1}_{\{S > 0\}}]$ as $x \uparrow b^*$. Meanwhile $G'(b^*) = \mathbb{E}[Q'(b^*+S)]$. The difference is $Q'(b^*) \mathbb{P}(S = 0)$, and $\mathbb{P}(S = 0) = 0$ exactly when 0 is regular for $(0, \infty)$ ([10]), which holds when $\sigma > 0$. $\square$

The two levels of the cascade are: (1) solve $q\mathcal{R}_q^+Q = G$ for Q ; (2) apply $q\mathcal{R}_q^+$ to Q^+ . Only the supremum factor enters, because the stopping region lies above the threshold. For a put-type payoff the infimum factor plays the same role ([8], Theorem 7 and Corollary 2). The two factors of Theorem 1 are thus used one at a time, each on its own side.

Remark 1 (Two-sided problems). When G is not monotone the stopping region need not be a half-line, the representation of Theorem 2 does not apply, and the fixed point of Lemma 2 couples both factors. This paper does not treat that case.

4.3 The Fredholm equation of the first level

Level (1) is the Fredholm integral equation of the first kind

$$\int_{s \geq 0} Q(x+s) \mathbb{P}(S \in ds) = G(x), \quad x \in \mathbb{R}, \quad (4.1)$$

with a convolution kernel on the half-line: the law of S , whose transform is $\sqrt{q}/\phi_q^+$. In transform terms, Q has the transform of G divided by Φ_q^+ . Exponentials are eigenfunctions, so the equation is solved exactly whenever G is a sum of exponentials, and also for every smooth G when Φ_q^+ is rational.

Theorem 3 (Explicit averaging functions). Assume A1 and $q > 0$. (i) If $G(x) = \sum_{k=1}^m c_k e^{r_k x}$ with $\Phi_q^+(r_k) < \infty$ for every k , then $Q(x) = \sum_{k=1}^m c_k e^{r_k x} / \Phi_q^+(r_k)$ solves $q\mathcal{R}_q^+Q = G$. In particular, for the call $G(x) = e^x - K$ with $\Phi_q^+(1) < \infty$, $Q(x) = e^x / \Phi_q^+(1) - K$ and $e^{b^*} = K \Phi_q^+(1)$ ([16]). (ii) Suppose S is the mixture $v_1 \text{Exp}(\alpha_1) + v_2 \text{Exp}(\alpha_2)$ with $0 < \alpha_1 < \eta_+ < \alpha_2$, so that $\Phi_q^+(r) = \alpha_1 \alpha_2 (\eta_+ - r) / (\eta_+ (\alpha_1 - r)(\alpha_2 - r))$ (the Kou case,



Proposition 1). Let $G \in C^2$ with G, G', G'' bounded by $C(1 + e^{\theta x})$ for some $\theta < \alpha_1$. Then

$$Q = \frac{1}{\alpha_1 \alpha_2} (\alpha_1 - D)(\alpha_2 - D) E_+ G, \quad E_+ G(x) := \eta_+ \int_0^\infty e^{-\eta_+ s} G(x + s) ds, \quad D = \frac{d}{dx},$$

solves $q \mathcal{R}_q^+ Q = G$.

Proof. $\mathbb{E}[e^{r(x+S)}] = \Phi_q^+(r) e^{rx}$, and the equation is linear. For the call, $\Phi_q^+(0) = 1$ and Q is strictly increasing with root $\log(K \Phi_q^+(1))$.

(ii) For $M \sim \text{Exp}(\alpha)$ and $K \in C^1$ growing slower than $e^{\alpha x}$, integration by parts gives $\mathbb{E}[(\alpha - D)K(x + M)] = \alpha K(x)$. Apply this to each component of the mixture, with $K = (\alpha_2 - D)E_+ G$ for the first and $K = (\alpha_1 - D)E_+ G$ for the second. This gives

$$\mathbb{E}[(\alpha_1 - D)(\alpha_2 - D)E_+ G(x + S)] = \alpha_1 \alpha_2 E_+ G(x) - (v_1 \alpha_1 + v_2 \alpha_2) (E_+ G)'(x).$$

The weights of Theorem 2 of [8], $v_1 = \alpha_2(\eta_+ - \alpha_1)/(\eta_+(\alpha_2 - \alpha_1))$ and $v_2 = \alpha_1(\alpha_2 - \eta_+)/(\eta_+(\alpha_2 - \alpha_1))$, give $v_1 \alpha_1 + v_2 \alpha_2 = \alpha_1 \alpha_2 / \eta_+$. The right side is therefore $\alpha_1 \alpha_2 (1 - D/\eta_+) E_+ G$. Another integration by parts gives $(E_+ G)' = \eta_+ (E_+ G - G)$, that is $(1 - D/\eta_+) E_+ G = G$. The growth bounds make every boundary term vanish. $\square$

In words, for Kou jumps the inverse of the supremum average is a second-order differential operator applied to an exponential smoothing. The first level of the cascade is then exact and costs $O(1)$ per point.

4.4 The optimal stopping boundary

Corollary 2 (Optimal stopping boundary). Under the hypotheses of Theorem 2, the free boundary is the single point $\partial \mathcal{D} = \{b^*\}$, the root of the averaging function. For $G(x) = \sum_k c_k e^{r_k x}$ it solves $\sum_k c_k e^{r_k b} / \Phi_q^+(r_k) = 0$.

Proof. Theorem 2(ii) and Theorem 3(i). $\square$

For the call payoff of §6, Corollary 2 is a single algebraic equation; it is the entry point to Theorem 4.

5. Computation and Its Checks

This section states the cost of the cascade on a grid and describes the two computations that check the closed forms of §6 independently. Neither check uses the roots of $\kappa(r) = q$ or the Wiener–Hopf factors.

5.1 Cost of the cascade

Remark 2 (Cost per level). On a uniform grid of N points, each level of the cascade costs $O(N \log N)$ operations when the law of S is evaluated by the fast Fourier transform: level (1) divides one transform by Φ_q^+ , and level (2) is one convolution with the law of S . A



forward and an inverse transform cost $O(N \log N)$ each, and the pointwise quotient and positive part cost $O(N)$. When Φ_q^+ is rational, both levels are explicit (Theorem 3 and Theorem 4) and cost $O(1)$ per point.

The general (non-rational) branch needs Φ_q^+ on the grid. It can be computed by a numerical Hilbert transform of $\log(q - \psi)$, which splits the logarithm of the symbol into its two half-plane parts (Algorithm 1, Appendix C). In this paper the cascade is evaluated only in the Kou case, where every step is exact.

5.2 Check 1: the variational inequality on a grid

The variational inequality of Definition 3 is solved on a uniform grid in $x = \log(P/I)$ over $[-8, x^* + 2]$, with step h . The diffusion and drift use central differences. The jump integral of the piecewise-linear interpolant of V is computed exactly for the exponential kernels: its upward and downward parts each obey a first-order recursion along the grid, so the linear system in (V, U, D) (the value and the two partial jump sums) is sparse and banded. The boundary conditions are $V = 0$ at and below the left end, and $V = g$ at and beyond the right end, with the jump mass that lands there integrated exactly. The stopping set is found by policy iteration ([17]), which converges in finitely many steps. Read off the grid, the threshold is the first stopping node. It is refined by fitting the smooth-fit profile $\sqrt{V - g} \propto (x^* - x)$ through the four continuation nodes next to the boundary. Algorithm 4 gives the steps.

5.3 Check 2: Monte Carlo of threshold rules

For a barrier c , the rule $\tau_c = \inf\{t : X_t \geq c\}$ has value $\mathbb{E}_x[e^{-q\tau_c}(e^{X_{\tau_c}} - 1)] = \mathbb{E}_x[(e^{X_{\tau_c}} - 1)\mathbb{1}_{\{\tau_c < \mathbf{e}(q)\}}]$ (in units of I), and this is simulated exactly. Events arrive at rate $\lambda + q$: a jump with probability $\lambda/(\lambda + q)$, otherwise the killing time $\mathbf{e}(q)$. Between events the diffusion increment is Gaussian, and its maximum given both endpoints is drawn from the Brownian-bridge law, $M = \frac{1}{2}(a + b' + \sqrt{(b' - a)^2 - 2\sigma^2\Delta t \log U})$ ([18]). A barrier crossed by the diffusion is hit exactly, at $X_\tau = c$; a barrier crossed by a jump is passed with the overshoot. Seven barriers around the closed-form threshold are run on common paths. Paths that fall more than 12 log-units below the lowest barrier are stopped, which biases each value by less than 1.2×10^{-6} (the probability that $S \geq 12$, times the largest payoff). Algorithm 5 gives the steps.

6. Real Options under Kou Jumps: Green-Capacity Timing

We now specialise the cascade to the Kou double-exponential jump-diffusion and apply it to the firm's green-capacity timing problem. The principal result is Theorem 4, the closed-form threshold, which needs both positive roots of the Kou equation.

6.1 The firm's investment problem

A firm holds a perpetual option to install irreversible green capacity at a one-time cost. We measure the installed project in carbon-price units: exercising when the carbon price



is P yields $P - I$, where I is the investment cost per unit of abatement value. With discounting at the firm's hurdle rate q , the firm's value given the current price $P_0 = P$ is the perpetual American call

$$V(P) = \sup_{\tau \in \mathcal{T}} \mathbb{E}[e^{-q\tau} (P_\tau - I)^+ | P_0 = P]. \quad (6.1)$$

Definition 5 (Green-capacity timing problem). The carbon price is $P_t = P_0 e^{X_t}$, where $X_t = bt + \sigma W_t + \sum_{j \leq N_t} Y_j$ is a Kou jump-diffusion ([19]): drift b , Gaussian volatility $\sigma > 0$, a Poisson process N_t of intensity $\lambda > 0$, and jumps Y_j that are $\text{Exp}(\eta_+)$ upwards with probability $p \in (0, 1)$ and $-\text{Exp}(\eta_-)$ downwards otherwise, with $\eta_+ > 1$ and $\eta_- > 0$. Its cumulant on $(-\eta_-, \eta_+)$ is

$$\kappa(r) = br + \frac{1}{2}\sigma^2 r^2 + \lambda \left(\frac{p\eta_+}{\eta_+ - r} + \frac{(1-p)\eta_-}{\eta_- + r} - 1 \right),$$

and $\psi(\xi) = \kappa(i\xi)$. The expected growth rate of the price is $g := \kappa(1)$, that is $\mathbb{E}[P_t] = P_0 e^{gt}$, and A2 requires $q > g$. The firm seeks $V(P)$ and the optimal investment threshold P^* .

6.2 The Kou factors

Proposition 1 (Kou Wiener–Hopf factors). Under Definition 5 and $q > 0$, the equation $\kappa(r) = q$, with κ continued rationally, has exactly four real roots $-\beta_2 < -\eta_- < -\beta_1 < 0 < \alpha_1 < \eta_+ < \alpha_2$, and the factors of Theorem 1 are

$$\phi_q^+(\xi) = \sqrt{q} \frac{\eta_+ (\alpha_1 - i\xi)(\alpha_2 - i\xi)}{\alpha_1 \alpha_2 (\eta_+ - i\xi)}, \quad \phi_q^-(\xi) = \sqrt{q} \frac{\eta_- (\beta_1 + i\xi)(\beta_2 + i\xi)}{\beta_1 \beta_2 (\eta_- + i\xi)}.$$

The supremum S is the mixture $v_1 \text{Exp}(\alpha_1) + v_2 \text{Exp}(\alpha_2)$ with $v_1 = \alpha_2(\eta_+ - \alpha_1)/(\eta_+(\alpha_2 - \alpha_1))$ and $v_2 = \alpha_1(\alpha_2 - \eta_+)/(\eta_+(\alpha_2 - \alpha_1))$.

Proof. The root location and the transforms $\Phi_q^+(r) = \alpha_1 \alpha_2 (\eta_+ - r)/(\eta_+(\alpha_1 - r)(\alpha_2 - r))$ and $\Phi_q^-(r) = \beta_1 \beta_2 (\eta_- + r)/(\eta_-(\beta_1 + r)(\beta_2 + r))$, with the mixture law of S , are Theorem 2 of [8] (see also [20]). By Theorem 1, $\phi_q^\pm(\xi) = \sqrt{q}/\Phi_q^\pm(i\xi)$. The zeros of ϕ_q^+ are at $\xi = -i\alpha_1, -i\alpha_2$ and its pole at $\xi = -i\eta_+$, all in the lower half-plane; those of ϕ_q^- , at $i\beta_1, i\beta_2$ and $i\eta_-$, lie in the upper half-plane. Appendix B (Lemma 8) checks the product directly. $\square$

6.3 Closed-form threshold and its dependence on the hurdle rate

Theorem 4 (The Kou investment threshold). Under Definition 5 with $q > \kappa(1)$, equivalently $\alpha_1 > 1$: (i) the optimal investment threshold is

$$P^* = I \Phi_q^+(1) = I \frac{\alpha_1 \alpha_2 (\eta_+ - 1)}{\eta_+ (\alpha_1 - 1)(\alpha_2 - 1)},$$



and $\tau^* = \inf\{t : P_t \geq P^*\}$ is optimal; (ii) the value is

$$V(P) = I \sum_{i=1}^2 \frac{v_i}{\alpha_i - 1} \left(\frac{P}{P^*}\right)^{\alpha_i} \quad (P < P^*), \quad V(P) = P - I \quad (P \geq P^*),$$

and both continuous and smooth fit hold at P^* ; (iii) P^* is strictly decreasing in q , with $P^* \rightarrow \infty$ as $q \downarrow \kappa(1)$ and $P^* \rightarrow I$ as $q \rightarrow \infty$, and its elasticity is

$$\frac{d \log P^*}{d \log q} = -q \sum_{i=1}^2 \frac{1}{\kappa'(\alpha_i) \alpha_i (\alpha_i - 1)};$$

(iv) as $\lambda \rightarrow 0$, $P^* \rightarrow I\beta/(\beta - 1)$, where β is the positive root of $b\beta + \frac{1}{2}\sigma^2\beta^2 = q$, which is the McDonald–Siegel threshold [2].

Proof. $-(ii)$ In $x = \log P$ the payoff is $G(x)^+$ with $G(x) = e^x - I$, non-decreasing. By Theorem 3(i), $Q(x) = e^x/\Phi_q^+(1) - I$, which is strictly increasing with root $b^* = \log(I\Phi_q^+(1))$. Kou has $\sigma > 0$ and up-jumps, so $\mathbb{P}(S > y) > 0$ for every y and $\mathbb{P}(S = 0) = 0$, and Theorem 2 gives the threshold, the optimality of τ^* and both fit conditions. For $x < b^*$, the mixture law of S (Proposition 1) gives

$$V = \sum_i v_i \int_{b^*-x}^{\infty} \alpha_i e^{-\alpha_i s} \left(\frac{e^{x+s}}{\Phi_q^+(1)} - I \right) ds = I \sum_i v_i e^{-\alpha_i(b^*-x)} \left(\frac{\alpha_i}{\alpha_i - 1} - 1 \right),$$

which is (ii).

(iii) $\mathbf{e}(q)$ has the law of $\mathbf{e}(1)/q$, and the running supremum is non-decreasing in time, so S decreases stochastically and strictly in q ; hence $\Phi_q^+(1) = \mathbb{E}[e^S]$ decreases strictly. As $q \downarrow \kappa(1)$, $\alpha_1 \downarrow 1$ and the factor $(\alpha_1 - 1)^{-1}$ diverges. As $q \rightarrow \infty$, $\alpha_1 \uparrow \eta_+$ and $\alpha_2 \rightarrow \infty$, and the fraction tends to 1. The elasticity follows from $d\alpha_i/dq = 1/\kappa'(\alpha_i)$ (implicit differentiation of $\kappa(\alpha_i) = q$) and $d \log \Phi_q^+(1)/d\alpha_i = -1/(\alpha_i(\alpha_i - 1))$.

(iv) As $\lambda \rightarrow 0$, $\kappa \rightarrow br + \frac{1}{2}\sigma^2 r^2$ away from the pole at η_+ . One of α_1, α_2 tends to β and the other to η_+ : the root beyond the pole collapses onto it when $\beta < \eta_+$, and the root below the pole does when $\beta > \eta_+$. In either case the formula in (i) tends to $\beta\eta_+(\eta_+ - 1)/(\eta_+(\beta - 1)(\eta_+ - 1)) = \beta/(\beta - 1)$. $\square$

Remark 3 (One exponential is not enough). The diffusion recipe of a single power, $V = AP^{\alpha_1}$ below the threshold with continuous and smooth fit, gives $P^* = I\alpha_1/(\alpha_1 - 1)$. Under up-jumps this is wrong: the equation $(q - \mathcal{L})V = 0$ below P^* sees V above P^* through the overshoot, and the three conditions (continuous fit, smooth fit, and the payoff paid on overshooting jumps) need both exponents. At the fit of §7 and $q = 0.10$ the one-root formula gives 3.2941 against 3.2783. The gap is small here because the fitted up-jumps are small; it is the same family of error that [8], §7, corrects for the put.

Theorem 4(iii) shows that, for any Lévy model, the threshold diverges as the hurdle rate falls to the expected growth rate of the price. A threshold that stays flat across hurdle rates would need the expected growth rate to sit far below the whole range of hurdle rates and the supremum law to barely move with the horizon. The data of §7 decide how close



the Kou threshold comes to this.




Figure 2: Value of the green-capacity option under the Kou model fitted in Section 7, with the expected price growth equal to the euro risk-free rate, for hurdle rates $q = 0.07, 0.10, 0.18$ (left; squares mark the thresholds of Theorem 4; teal dots are the grid solution of the variational inequality at $q = 0.10$, step $h = 0.00125$), and the option premium $V - (P - I)^+$ (right), which vanishes from the threshold on. Prices and values are in units of the investment cost I .

7. EU ETS Calibration and the Threshold

We fit the Kou jump-diffusion to EU ETS primary-auction prices, apply Theorem 4, and compare the threshold with the McDonald–Siegel threshold of a geometric Brownian motion fitted to the same prices.

7.1 EU ETS auction data

The price data are the clearing prices of the EU ETS primary auctions held by the European Energy Exchange (EEX), as compiled and published by the Intrat Foundation’s open-data service [3] under the Creative Commons Attribution 4.0 licence. The licence was read before use; it requires attribution, which this paragraph, Figure 3 and the reference list give. Every auction from January 2018 up to and including 21 May 2025, the date of this paper, was retrieved: 1,567 auctions, the first on 2018-01-08 and the last on 2025-05-20. Nothing later is used. The file was checked against EEX’s own auction reports: each of the 1,567 prices matches an EEX auction of the same date to the cent. The raw file, the query, the retrieval time and its checksum stay with the paper’s computation package; the paper publishes only estimates and the chart below.

We use the log-returns between consecutive auctions, 1,566 in all, and measure the time between auctions in business days ($\Delta = \text{business days}/252$). Of the gaps, 1,360 are one business day; the longest are holiday breaks of up to 34 business days. The 1,360 one-day returns have skewness -0.03 and excess kurtosis 2.83. The December 2024 auctions (10 of



them) cleared at 67.22 €/tCO₂ on average, and the last auction before this paper's date, on 2025-05-20, cleared at 70.69 €/tCO₂.




Figure 3: EU ETS primary-auction clearing prices, 2018-01-08 to 2025-05-20 (left), and the log-returns between consecutive auctions (right; teal dots: the two moves beyond 15 % in absolute value; dashed: ± 15 %). The teal dot on the left marks the last auction before the paper's date. Data: EEX auction results, compiled and published by the Infracore Foundation (energy.infracore.pl), CC BY 4.0; returns computed by us.

7.2 Kou maximum-likelihood fit and its checks

The return over a gap of length Δ has the exact Kou transition density $f_{\Delta}(x) = (2\pi)^{-1} \int e^{-iux} e^{\Delta\psi(u)} du$. We compute it by the fast Fourier transform on a grid of 8192 points over $[-1.5, 1.5]$, once for each distinct gap, and interpolate it at the observed returns. The log-likelihood $\ell(\theta) = \sum_k \log f_{\Delta_k}(x_k; \theta)$, with $\theta = (b, \sigma, \lambda, p, \eta_+, \eta_-)$, is maximised by L-BFGS-B from 18 starting points in transformed coordinates, followed by a Nelder–Mead polish of the best. Standard errors come from the inverse of the numerical Hessian of $-\ell$ (Algorithm 3). For comparison, a geometric Brownian motion is fitted by its closed-form maximum likelihood on the same gaps.



Table 1: Maximum-likelihood fits to 1,566 EU ETS auction-to-auction log-returns, 2018-01-08 to 2025-05-20. Kou: exact transition density; standard errors from the observed information; mean and standard deviation over 500 parametric-bootstrap refits. GBM: closed-form estimates, standard errors $s/\sqrt{T}$ and $s/\sqrt{2n}$ ($T = 7.62$ business years).

Parameter	Estimate	Standard error	Bootstrap mean	Bootstrap s.d.
b (drift of $\log P$, yr^{-1})	0.035	0.325	−0.071	0.650
σ ($\text{yr}^{-1/2}$)	0.345	0.015	0.339	0.026
λ (jumps per year)	54.3	26.3	79.2	103.9
p (probability of an up-jump)	0.70	0.14	0.69	0.15
η_+ (up-jump rate)	46.1	11.7	49.9	16.3
η_- (down-jump rate)	28.4	7.4	30.8	9.9
GBM: m (drift of $\log P$, yr^{-1})	0.289	0.159	—	—
GBM: s ($\text{yr}^{-1/2}$)	0.440	0.008	—	—
Log-likelihood: Kou / GBM	3383.5 / 3309.5			

The fit puts 39 % of the return variance into jumps. Up-jumps are more frequent ($p = 0.70$) but smaller: mean 2.2 % against 3.5 % for down-jumps. The jump intensity is the least precise parameter (standard error 26.3 on 54.3 per year). The Kou log-likelihood exceeds the Gaussian one by 74.0, and the AIC difference is 140.0 in favour of Kou.

The fit is checked in four ways.

- (a) The Fourier density is checked against 10^6 increments simulated exactly from the fitted model, at gaps of 1, 3 and 34 business days. The Kolmogorov–Smirnov distances are 1.4×10^{-3} , 8.8×10^{-4} and 7.9×10^{-4} .
- (b) A parametric bootstrap draws 500 data sets exactly from the fitted model on the observed gaps and refits each one (Table 1). The bootstrap means are close to the estimates, except for λ : its sampling distribution is skewed to the right (bootstrap median 62.0, mean 79.2). The bootstrap standard deviations are 1.0–3.9 times the Hessian standard errors, the largest ratio being for λ . The observed information therefore understates the uncertainty of the jump parameters, and the threshold bands of §7.3 use the bootstrap.
- (c) Under each fitted model, the probability integral transforms of the 1,566 returns



are tested for uniformity (Figure 4, right). The Kou fit passes ($KS = 0.0116$, $p = 0.98$); the Gaussian fit is rejected ($KS = 0.0433$, $p = 0.0054$).

- (d) The model's one-day skewness and excess kurtosis are -0.36 and 5.31, against -0.03 and 2.83 in the sample. The model's tails are heavier than the one-day sample moments suggest; the multi-day gaps, which the likelihood also fits, carry part of this.




Figure 4: Fit of the Kou model to the EU ETS auction returns. Left: the density of one-business-day returns (teal dots, histogram on a log scale) against the fitted Kou density (red) and the fitted Gaussian (dashed). Right: probability integral transforms of all returns under each fitted model, as quantiles minus the uniform quantile, with the 5 % Kolmogorov–Smirnov band (shaded, teal edges).

The drift. Under the fitted model, the sample growth rate of the price is $\hat{\kappa}(1) = 0.386$ per year, with standard error 0.159. This is above every hurdle rate considered below, and it is statistically almost uninformative: seven years of prices cannot pin down a growth rate. Used as the expected growth rate, it would make waiting forever optimal at every hurdle rate (Theorem 4(iii)). We therefore take the drift from an economic argument rather than from the sample, as is standard in real-options valuation [1]. An allowance is a bankable asset with no yield, so cost of carry ties its forward price to the spot by the risk-free rate, and the market's expected growth of the price is that rate. We set $g = \kappa(1) = r$, the euro-area AAA 10-year spot rate of the European Central Bank on 2025-05-21, $r = 2.7335\%$ [21]. The volatility and jump parameters are those of Table 1, and the drift b is set so that $\kappa(1) = r$.

7.3 The threshold against McDonald–Siegel

Table 2 gives the Kou threshold of Theorem 4 at the hurdle rates 0.07, 0.10, 0.12, 0.15 and 0.18, the 95 % bootstrap band, and the McDonald–Siegel threshold of the geometric Brownian motion fitted to the same prices, with the same expected growth rate.



Table 2: The investment threshold of Theorem 4 at the Kou fit of Table 1, with expected price growth $g = r = 2.7335$ %, against McDonald–Siegel for the GBM fit with the same growth. ‘Difference’: McDonald–Siegel relative to Kou. ‘Elasticity’: $d \log P^*/d \log q$ by Theorem 4(iii) (a finite difference agrees to 4×10^{-9}). ‘Trigger cost’: the largest investment cost (€/tCO₂) whose threshold is met at the auction price of 70.69 €/tCO₂ on 2025-05-20. Bootstrap intervals from 500 refits.

q	α_1	α_2	Kou P^*/I	95 % bootstrap	McDonald–Siegel P^*/I	difference	elasticity	trigger cost
0.07	1.2812	58.80	4.5344	4.231–4.840	4.5442	+0.22 %	−1.086	15.6
0.10	1.4359	58.81	3.2783	3.091–3.466	3.2871	+0.27 %	−0.765	21.6
0.12	1.5277	58.81	2.8810	2.729–3.033	2.8894	+0.29 %	−0.657	24.5
0.15	1.6534	58.81	2.5184	2.399–2.637	2.5265	+0.32 %	−0.553	28.1
0.18	1.7679	58.81	2.2913	2.192–2.390	2.2991	+0.34 %	−0.486	30.9

The Kou threshold multiple falls from 4.534 at $q = 0.07$ to 2.291 at $q = 0.18$, by a factor of 1.98. The McDonald–Siegel multiple falls from 4.544 to 2.299, and it lies 0.22–0.34 % above the Kou multiple at every rate. Both sit inside the Kou bootstrap band, and in 15 % of the bootstrap refits (each refitting the Kou and the Gaussian model to the same simulated prices) the order of the two thresholds is reversed. The elasticity of the threshold with respect to the hurdle rate is between 1.09 and 0.49 in absolute value under Kou, against 1.08 to 0.48 under McDonald–Siegel. On these data, then, the jump-aware threshold is not stable across hurdle rates, and the McDonald–Siegel threshold is not unstable relative to it. The two agree to within half a per cent.

The reason is the horizon. The threshold depends on the law of the supremum over an exponential time with mean $1/q$, that is 5 to 15 years. Over such horizons a jump-diffusion with 54 small jumps a year is close to a Brownian motion with the same variance. The jumps’ asymmetry and the overshoot (Remark 3) move the threshold by less than the sampling error. A model that keeps only the diffusion part of the fit, with volatility $\hat{\sigma} = 0.345$, would understate the threshold multiple: 2.683 at $q = 0.10$ against 3.278. The jumps matter through the variance they carry, not through their shape.

Table 3 gives the two independent checks.



Table 3: Theorem 2, Theorem 3 and Theorem 4 checked two ways in the call case. Grid: policy-iteration solve of the variational inequality; its stopping set is a single half-line at every step size, as Theorem 2 states, and policy iteration started from no stopping and from stopping wherever $g > 0$ reaches the same stopping set and grid values within 1×10^{-300} at $h = 0.0025$ (Theorem 5 and Corollary 1, discretely). Further grid details (Section 5.2): threshold at $h = 0.00125$ refined by the smooth-fit profile (Theorem 4: Table 2), and the largest error of the grid value against Theorem 4(ii) over $\log(P/I) \in [-3, \log(P^*/I) + 1]$ at $h = 0.00125$; at $h = 0.005$ and 0.0025 the errors are 4.0–4.0 times larger per halving. Monte Carlo: exact simulation of the rule τ^* from $P = I$ (Section 5.3), 1,000,000 paths, $\pm$ one standard error; z is the Monte Carlo value minus Theorem 4, in standard errors.

q	grid P^*/I	value error	$V(I)/I$, Theorem 4	$V(I)/I$, Monte Carlo	z
0.07	4.5347	6.6×10^{-5}	0.5095	0.5084 ± 0.0012	−0.93
0.10	3.2787	4.1×10^{-5}	0.4142	0.4141 ± 0.0009	−0.10
0.12	2.8807	3.3×10^{-5}	0.3735	0.3741 ± 0.0008	+0.82
0.15	2.5186	2.6×10^{-5}	0.3297	0.3299 ± 0.0006	+0.25
0.18	2.2913	2.2×10^{-5}	0.2981	0.2979 ± 0.0005	−0.42

The grid solve of the variational inequality converges at second order: the largest value error falls by a factor of about 4.0–4.0 each time h halves. At $h = 0.00125$ its threshold agrees with Theorem 4 to within 1.3×10^{-4} relative, and the stopping set is a single half-line at every step size. The Monte Carlo value of the rule τ^* agrees with the closed-form value $V(I)$ within 0.9 standard errors at every hurdle rate. On common paths, moving the barrier by ± 0.2 in $\log P$ lowers the value at every hurdle rate, by at least 4.4 standard errors of the difference. Moving it by ± 0.1 lowers the value at every hurdle rate from 10 % (by at least 2.6 standard errors). The smaller moves change the value by at most 2.4 standard errors, because the value is flat at its maximum. No barrier does better than τ^* by more than 0.6 standard errors (Figure 5, right).




Figure 5: Investment threshold multiple P^*/I against the hurdle rate (left): Theorem 4 for the fitted Kou model (red), McDonald–Siegel for the fitted geometric Brownian motion (dashed), the grid solution of the variational inequality (teal dots), and 95 % intervals from 500 parametric-bootstrap refits; the expected price growth is the euro risk-free rate. Right: Monte Carlo values of threshold rules τ_c with barriers around $\log(P^*/I)$, relative to the closed-form value of Theorem 4 at $P = I$, with 95 % error bars (10^6 paths per hurdle rate, common paths across barriers).

7.4 Sensitivity

Table 4 repeats the comparison under two changes. In the first, the expected growth rate is zero instead of the risk-free rate. In the second, the model is refitted to the Phase 4 sample only (946 returns from 2021-01-29).

Table 4: Threshold multiples P^*/I (Kou: Theorem 4; MS: McDonald–Siegel) for the main fit with $g = r$, the main fit with $g = 0$, and a refit to the Phase 4 sample (2021-01-29 to 2025-05-20; $\hat{\sigma} = 0.332$, $\hat{\lambda} = 36.2$, $\hat{p} = 0.71$, $\hat{\eta}_+ = 41.3$, $\hat{\eta}_- = 27.8$; GBM $s = 0.408$) with $g = r$.

q	Kou, main	MS, main	Kou, $g = 0$	MS, $g = 0$	Kou, Phase 4	MS, Phase 4
0.07	4.534	4.544	3.045	3.053	4.189	4.196
0.10	3.278	3.287	2.571	2.578	3.064	3.071
0.12	2.881	2.889	2.378	2.386	2.707	2.714
0.15	2.518	2.526	2.179	2.187	2.381	2.388
0.18	2.291	2.299	2.041	2.049	2.176	2.183

Both changes move the level of the threshold, and neither changes the comparison: Kou and McDonald–Siegel stay within 0.4 % of each other, McDonald–Siegel above. The



dependence on the hurdle rate remains strong. From 7 % to 18 % the threshold multiple falls by a factor of 1.49 with zero expected growth, and by 1.92 with the Phase 4 refit (against 1.98 in the main case).

7.5 Implications

Theorem 4 is homogeneous of degree one in the investment cost, so the data give a threshold *multiple*. At the last auction price before this paper's date, 70.69 €/tCO₂, the investment trigger is met only by projects whose cost per unit of abatement value is at most 70.69/(P^*/I). That is 15.6 €/tCO₂ at a 7 % hurdle rate and 30.9 €/tCO₂ at 18 % (Table 2). Three implications follow, within the model and its assumptions. First, the hurdle rate and the expected growth rate of the carbon price are the inputs that matter most: between them they move the trigger by a factor of two or more. Second, at the fitted jump sizes, modelling jumps does not change the trigger by more than a fraction of a per cent once the variance is fitted to the same data; the McDonald–Siegel formula remains a good approximation for this payoff. Third, omitting the jump variance, by fitting only the diffusion part, understates the trigger. None of these is a recommendation to any firm. Project costs, the expected growth rate of the price, and the horizon of the policy regime are outside the model.

8. Conclusion

The paper has developed a Wiener–Hopf cascade for optimal stopping under Lévy noise and applied it to the timing of irreversible green-capacity investment under a Kou jump-diffusion fitted to EU ETS auction prices.

8.1 Summary of results

Claim 1. Theorem 1 factorises the resolvent symbol as $q - \psi = \phi_q^+ \phi_q^-$, with $\phi_q^\pm$ the reciprocal transforms of the supremum and the infimum at an exponential time. It rests on the factorisation of the published resolvent paper [8].

Claim 2. Theorem 2 solves every one-sided problem in two levels: an averaging function from a first-kind Fredholm equation with the law of the supremum as kernel, and the supremum-average of its positive part. Theorem 3 makes the first level explicit for exponential-type payoffs and for rational factors. Corollary 2 identifies the boundary as the root of the averaging function.

Claim 3. Theorem 4 gives the Kou investment threshold $P^* = I \Phi_q^+(1)$, with two roots, smooth fit, strict monotonicity in the hurdle rate, divergence as the hurdle rate falls to the price's growth rate, and the McDonald–Siegel limit as jumps vanish.

Claim 4. An exact-likelihood Kou fit to 1,566 EU ETS auction returns passes a probability-integral-transform test that rejects the Gaussian model; jumps carry 39 % of the variance.

Claim 5. With the price expected to grow at the risk-free rate, the threshold multiple



falls from 4.53 to 2.29 as the hurdle rate rises from 7 % to 18 %, and the McDonald–Siegel multiple is within half a per cent of it throughout. The grid solve and the Monte Carlo confirm the closed form.

8.2 Limits and extensions

The problem is perpetual and the parameters are constant. A finite investment window turns the threshold into a time-dependent boundary, and the cascade then has to be run on each time slice of a backward recursion; this paper does not study that scheme's convergence. The expected growth rate of the price is an assumption, not an estimate, and Theorem 4(iii) shows how strongly the threshold depends on it. The jump intensity is imprecisely estimated, but the threshold barely depends on it (Table 2). Two-sided stopping problems (Remark 1) and stopping games under jump-driven processes need both factors at once and are not treated here.

A. Comparison Principle: Full Proof

This appendix proves Lemma 3 by doubling of variables adapted to the non-local generator. The argument combines the technique of Crandall–Ishii–Lions [11] with the Jakobsen–Karlsen non-local maximum principle [12] and the small-jump/large-jump split of Barles–Imbert [13]; [8], Theorem 5, gives the same argument for payoffs of linear growth.

A.1 Setup and localisation

Assumption 3 (Hypotheses recalled). $\underline{V}$ is an upper-semicontinuous viscosity subsolution and $\overline{V}$ a lower-semicontinuous viscosity supersolution of the variational inequality (Def. 3), both in $\mathcal{G}$, and A1–A2 hold. In the convention $\min(qV - \mathcal{L}V, V - g) = 0$, a subsolution satisfies $\min(\cdot) \leq 0$ and a supersolution $\min(\cdot) \geq 0$ in the viscosity sense; in particular $\overline{V} \geq g$.

Lemma 4 (Localisation). Since $\kappa(1) < q$ and $\kappa(0) = 0 < q$, there are $\theta \in (1, \eta_+)$ and $\theta' \in (0, \eta_-)$ with $\kappa(\theta) < q$ and $\kappa(-\theta') < q$. Let $w(x) := e^{\theta x} + e^{-\theta' x}$. Then $(q - \mathcal{L})w > 0$, and for every $\varepsilon > 0$, $\underline{V}_\varepsilon := \underline{V} - \varepsilon w$ is a subsolution with $\underline{V}_\varepsilon - \overline{V} \rightarrow -\infty$ as $|x| \rightarrow \infty$.

Proof. $(q - \mathcal{L})w = (q - \kappa(\theta))e^{\theta x} + (q - \kappa(-\theta'))e^{-\theta' x} > 0$ by Definition 2. Subtracting εw lowers $qV - \mathcal{L}V$ and $V - g$, so the subsolution inequality is preserved. Both functions are in $\mathcal{G}$, which grows at most like e^x at $+\infty$ and stays bounded at $-\infty$, so the weight dominates at both ends. $\square$

Suppose, for contradiction, that $M_0 := \sup_{\mathbb{R}}(\underline{V} - \overline{V}) > 0$. Fix $\varepsilon > 0$ small enough that $M_\varepsilon := \sup(\underline{V}_\varepsilon - \overline{V}) > M_0/2$. By Lemma 4 the supremum is attained on a compact set K^* . For $\eta > 0$ define on $\mathbb{R}^2$

$$\Phi_\eta(x, y) := \underline{V}_\varepsilon(x) - \overline{V}(y) - \frac{(x - y)^2}{2\eta}. \quad (\text{A.1})$$



Lemma 5 (Maximum approaches the diagonal). For η small, Φ_η attains its maximum $M_\eta \geq M_\varepsilon$ at some (x_η, y_η) in a fixed compact set, and $(x_\eta - y_\eta)^2/\eta \rightarrow 0$ as $\eta \rightarrow 0^+$, with limit points $(\hat{x}, \hat{x})$ at which $\underline{V}_\varepsilon - \bar{V}$ attains M_ε .

Proof. The standard penalisation argument of [11], §3, using the coercivity of Lemma 4. $\square$

A.2 Non-local maximum principle

Lemma 6 (Jakobsen–Karlsen non-local maximum principle). If $\underline{V}_\varepsilon(x_\eta) > g(x_\eta)$, then, with $p_\eta = (x_\eta - y_\eta)/\eta$, there are $X_\eta \leq Y_\eta$ (the second-order parts of the semijets) and, for every $\delta \in (0, 1)$, non-local terms $J_{\eta,\delta}^\pm$ such that

$$q \underline{V}_\varepsilon(x_\eta) - b p_\eta - \frac{1}{2} \sigma^2 X_\eta - J_{\eta,\delta}^+ \leq 0, \quad q \bar{V}(y_\eta) - b p_\eta - \frac{1}{2} \sigma^2 Y_\eta - J_{\eta,\delta}^- \geq 0.$$

Here $J_{\eta,\delta}^\pm$ split into a small-jump part over $|z| \leq \delta$, evaluated on the quadratic test function, and a large-jump part over $|z| > \delta$, evaluated on $\underline{V}_\varepsilon$ and $\bar{V}$. The small-jump parts satisfy $|J_{\eta,\delta}^{+, \text{small}} - J_{\eta,\delta}^{-, \text{small}}| \leq \eta^{-1} \int_{|z| \leq \delta} z^2 \nu(dz)$, and the large-jump parts satisfy

$$J_{\eta,\delta}^{+, \text{large}} - J_{\eta,\delta}^{-, \text{large}} = \int_{|z| > \delta} \left[(\underline{V}_\varepsilon(x_\eta + z) - \bar{V}(y_\eta + z)) - (\underline{V}_\varepsilon(x_\eta) - \bar{V}(y_\eta)) \right] \nu(dz) \leq 0.$$

Proof. The jets are Theorem 3.1 of [12], with the split of [13], §2; $X_\eta \leq Y_\eta$ is the matrix inequality of the Crandall–Ishii lemma in one dimension. The large-jump bound holds because $\Phi_\eta(x_\eta + z, y_\eta + z) \leq \Phi_\eta(x_\eta, y_\eta)$: a common shift leaves the penalty unchanged. $\square$

A.3 Contradiction and conclusion

If $\underline{V}_\varepsilon(x_\eta) \leq g(x_\eta)$ along a subsequence, then, since $\bar{V} \geq g$, $M_\varepsilon \leq M_\eta \leq g(x_\eta) - g(y_\eta) \rightarrow 0$ by the uniform continuity of g on compacts, which contradicts $M_\varepsilon > 0$. Otherwise Lemma 6 applies. Subtracting its two inequalities, the drift terms cancel, the second-order terms contribute $\frac{1}{2} \sigma^2 (X_\eta - Y_\eta) \leq 0$, and the large-jump terms contribute at most 0:

$$q M_\eta \leq q (\underline{V}_\varepsilon(x_\eta) - \bar{V}(y_\eta)) \leq \eta^{-1} \int_{|z| \leq \delta} z^2 \nu(dz). \quad (\text{A.2})$$

Let $\delta \rightarrow 0$ for fixed η , so the right side tends to 0 while $J_{\eta,\delta}^\pm$ remain admissible. This gives $q M_\varepsilon \leq q M_\eta \leq 0$, a contradiction. Hence $\underline{V} - \varepsilon w \leq \bar{V}$ for every $\varepsilon > 0$, and letting $\varepsilon \rightarrow 0$ gives $\underline{V} \leq \bar{V}$.

Theorem 5 (Comparison principle, restated). Under Assumption 3, $\underline{V} \leq \bar{V}$ on $\mathbb{R}$.

Proof. Lemmas 4, 5 and 6 and the argument above. $\square$

Remark 4 (The growth class). The weight w needs $q > \kappa(1)$, which is A2: without it, the value of the call is infinite and the class $\mathcal{G}$ contains no solution. For payoffs bounded at $+\infty$ the weight $e^{\theta x}$ may be taken with any $\theta \in (0, \eta_+)$ such that $\kappa(\theta) < q$.



B. Explicit Kou Wiener-Hopf Factors

This appendix gives the root structure and the direct verification of the factors of Proposition 1, with Definition 5's notation.

B.1 The four roots of the Kou equation

Clearing denominators in $\kappa(r) = q$ gives the quartic

$$\begin{aligned} D(r) &:= (q - \kappa(r))(\eta_+ - r)(\eta_- + r) \\ &= (q + \lambda - br - \tfrac{1}{2}\sigma^2 r^2)(\eta_+ - r)(\eta_- + r) - \lambda p \eta_+(\eta_- + r) - \lambda(1 - p) \eta_-(\eta_+ - r), \end{aligned} \tag{B.1}$$

with leading coefficient $\frac{1}{2}\sigma^2 > 0$.

Lemma 7 (Root location). Under Definition 5 and $q > 0$, D has four real roots, $-\beta_2 < -\eta_- < -\beta_1 < 0 < \alpha_1 < \eta_+ < \alpha_2$, and $D(r) = \frac{1}{2}\sigma^2(\alpha_1 - r)(\alpha_2 - r)(\beta_1 + r)(\beta_2 + r)$.

Proof. On $(0, \eta_+)$, $\kappa - q$ runs from $-q$ to $+\infty$ (the pole at η_+); on (η_+, ∞) it runs from $-\infty$ to $+\infty$, since $\sigma > 0$. On the negative side the same holds about $-\eta_-$. This gives four sign changes of $q - \kappa$ in the stated intervals, which are roots of D because $(\eta_+ - r)(\eta_- + r) \neq 0$ there. A quartic has no more roots. This is [8], Theorem 2. $\square$

B.2 Rational form of the factors

Lemma 8 (Kou factor formulas). With the roots of Lemma 7, the functions $\phi_q^\pm$ of Proposition 1 satisfy $\phi_q^+(\xi) \phi_q^-(\xi) = q - \psi(\xi)$ on the strip $-\eta_+ < \Im \xi < \eta_-$, and $\phi_q^\pm(0) = \sqrt{q}$.

Proof. At $r = 0$, Lemma 7 reads $q \eta_+ \eta_- = \frac{1}{2}\sigma^2 \alpha_1 \alpha_2 \beta_1 \beta_2$. With $r = i\xi$,

$$\phi_q^+ \phi_q^- = q \frac{\eta_+ \eta_-}{\alpha_1 \alpha_2 \beta_1 \beta_2} \frac{(\alpha_1 - r)(\alpha_2 - r)(\beta_1 + r)(\beta_2 + r)}{(\eta_+ - r)(\eta_- + r)}$$

and, by Lemma 7, the right side equals $D(r)/((\eta_+ - r)(\eta_- + r)) = q - \kappa(r) = q - \psi(\xi)$. The normalisation is immediate at $\xi = 0$. $\square$

B.3 Numerical verification

The script that produces every number of the paper computes the four roots in two ways: as eigenvalues of the companion matrix of D , and by bracketing each sign change of $\kappa - q$ with Brent's method. It then evaluates the factor product of Proposition 1 and Lemma 8 against the symbol, and the modulus of ϕ_q^+ on the upper half-plane (Table 5). As $\lambda \rightarrow 0$ ($\lambda = 10^{-9}$), the threshold of Theorem 4 matches McDonald–Siegel with the diffusion volatility to within 5.1×10^{-12} (Theorem 4(iv)).



Table 5: Theorem 1, Proposition 1 and Lemma 8 checked at the EU ETS fit (drift set by $\kappa(1) = r$). The four roots of the Kou equation by companion-matrix eigenvalues against Brent's method on each bracket of Lemma 7 (largest relative gap); the factor product against the symbol at 2 000 random points of the strip $-\eta_+ < \Im \xi < \eta_-$, $|\Re \xi| \leq 200$ (largest relative difference); and the smallest modulus of ϕ_q^+ at 500 random points of the upper half-plane, which stays away from zero.

q	α_1	α_2	β_1	β_2	root gap	product error	$\min \phi_q^+ $
0.07	1.2812	58.80	0.5604	36.39	2.1×10^{-14}	1.6×10^{-13}	1.00
0.10	1.4359	58.81	0.7142	36.39	4.2×10^{-14}	3.0×10^{-14}	1.06
0.12	1.5277	58.81	0.8054	36.39	1.4×10^{-14}	3.4×10^{-14}	0.85
0.15	1.6534	58.81	0.9302	36.40	1.8×10^{-14}	4.0×10^{-14}	1.70
0.18	1.7679	58.81	1.0438	36.40	7.2×10^{-15}	3.7×10^{-14}	2.71

C. Algorithms

This appendix gives the routines behind the numbers of §7: the general cascade (stated for completeness; only its rational branch is run here), the Kou threshold, the Kou fit, and the two checks.

C.1 Cascade construction

Algorithm 1 — Two-level Wiener–Hopf cascade for a one-sided problem

Computes the averaging function, the threshold and the value for a non-decreasing G , a Lévy exponent and a discount rate.

```

1  Input:  exponent psi(xi); discount q; non-decreasing G(x); grid x[1..N],
         step h.
2  Output: averaging function Q[1..N]; threshold b*; value V[1..N].
3
4  1. Supremum factor Phi+(i xi) = E[exp(i xi S)]:
5     if the exponent is rational (Kou, hyperexponential):
6         roots of kappa(r) = q (Algorithm 2) -> Phi+ in closed form
7     else:
8         L(xi) = log(q - psi(xi)) on the dual grid (damped if G grows);
9         log phi+ = upper-half-plane part of L by a discrete Hilbert
10        transform (FFT); Phi+ = sqrt(q) / phi+.
11  2. Level 1 (Fredholm equation of the first kind):
12     Q = inverse transform of [ transform(G) / Phi+ ] (Theorem 3)
13  3. b* = the root of Q (bisection on the grid, then linear interpolation).
14  4. Level 2: V = E[ Q(x + S)^+ ] = inverse transform of
15         [ transform(max(Q, 0)) * Phi+ ]
16  5. return Q, b*, V

```



Each step is one or two transforms of length N : $O(N \log N)$ (Remark 2).

C.2 Kou threshold and value

Algorithm 2 — Kou threshold, Theorem 4

Evaluates the roots, the threshold multiple and the value function from the Kou parameters and the discount rate.

```

1 Input:  sigma, lambda, p, eta_up, eta_dn; growth  $g = \kappa(1)$ ; discount  $q > g$ .
2 Output: alpha1, alpha2, beta1, beta2;  $P^*/I$ ;  $V(P)/I$ .
3
4 1.  $b = g - \sigma^2/2 - \lambda(p\eta_{up}/(\eta_{up}-1) + (1-p)\eta_{dn}/(\eta_{dn}+1) - 1)$ 
5 2.  $D(r) = (q + \lambda - b r - \sigma^2 r^2/2)(\eta_{up} - r)(\eta_{dn} + r) - \lambda p \eta_{up} (\eta_{dn} + r) - \lambda (1-p) \eta_{dn} (\eta_{up} - r)$ 
6
7 roots of  $D$ : companion-matrix eigenvalues; check each against Brent's
8 method on its bracket (Lemma 7).
9 3.  $\text{ratio} = \alpha_1 \alpha_2 (\eta_{up} - 1) / (\eta_{up} (\alpha_1 - 1)(\alpha_2 - 1))$ 
10 4.  $v_1 = \alpha_2 (\eta_{up} - \alpha_1) / (\eta_{up} (\alpha_2 - \alpha_1))$ 
11     $v_2 = \alpha_1 (\alpha_2 - \eta_{up}) / (\eta_{up} (\alpha_2 - \alpha_1))$ 
12 5.  $V(P)/I = \sum_i v_i / (\alpha_i - 1) (P/(\text{ratio } I))^{\alpha_i}$  for  $P < \text{ratio } I$ ,
13     $= P/I - 1$  otherwise.
```

C.3 EU ETS Kou calibration

Algorithm 3 — Exact-likelihood Kou fit

Maximum-likelihood estimation of the Kou parameters from returns observed over unequal gaps, with observed-information standard errors.

```

1 Input:  returns  $x[1..n]$ ; gaps  $\Delta[1..n]$  (business days / 252).
2 Output: theta = (b, sigma, lambda, p, eta_up, eta_dn); standard errors.
3
4 1. Dual grid  $u_k = 2 \pi k / (M \Delta x)$ ,  $M = 8192$  points over  $[-1.5, 1.5]$ .
5 2. For theta:  $\psi_k = \psi(u_k; \theta)$ ; for each distinct gap  $\Delta$ :
6     $f_{\Delta} = \text{Re FFT}[ \exp(\Delta \psi_k) * \exp(-i u_k x_0) ] / (M \Delta x)$ 
7 3.  $\text{loglik}(\theta) = \sum_j \log \text{interp}(x_j; f_{\{\Delta_j\}})$ 
8 4. Maximise in  $z = (b, \log \sigma, \log \lambda, \text{logit } p, \log(\eta_{up} - 1), \log \eta_{dn})$  by L-BFGS-B from 18 starts; polish the best by
9    Nelder-Mead and L-BFGS-B.
10
11 5. Standard errors: inverse of the central-difference Hessian of  $-\text{loglik}$ 
12    in theta.
```



C.4 The two checks

Algorithm 4 — Variational inequality on a log-price grid

Policy iteration for $\min((q - \mathcal{L})V, V - g) = 0$ with exact jump integrals for exponential kernels.

```

1 Input: Kou parameters with drift b; q; step h; x in [-8, x* + 2].
2 Output: grid value V; stopping set; refined threshold.
3
4 1. Up-jump weights: w_k = Cu exp(-a k), k >= 1, a = eta_up h,
5                   Cu = p eta_up h^2 (cosh a - 1)/a^2; likewise down.
6 2. Partial sums U_i = sum_{k>=1} w_k V_{i+k} and D_i (down) obey
7   U_i = exp(-a) (Cu V_{i+1} + U_{i+1}), D_i = exp(-c) (Cd V_{i-1} +
8   D_{i-1}),
9   with V = g at and beyond the right end (summed in closed form)
10  and V = 0 at and below the left end.
11 3. Continuation rows: (q + lambda + sigma^2/h^2 - lambda w0) V_i
12   - lo V_{i-1} - up V_{i+1} - lambda (U_i + D_i) = 0 (central
13   differences).
14   Stopping rows: V_i = g_i.
15 4. Policy iteration: start with no stopping; solve the sparse system;
16   stop at i where V_i - g_i < residual of the continuation row; repeat
17   until the stopping set is unchanged.
18 5. Threshold: first stopping node; refined by a line through
19   sqrt(V - g) at the four continuation nodes next to it.
```

Algorithm 5 — Monte Carlo of threshold rules

Exact simulation of $\mathbb{E}[(e^{X_\tau} - 1)\mathbb{1}_{\{\tau < e(q)\}}]$ for several barriers on common paths.

```

1 Input: Kou parameters with drift b; q; barriers c[1..K]; start x0 = 0; n
2   paths.
3 Output: values and standard errors for each barrier; covariances.
4
5 1. For each path: x = x0, running max m = x0, payoff[1..K] = 0.
6 2. Repeat until killed or every barrier is passed:
7   dt ~ Exp(lambda + q); x' = x + b dt + sigma sqrt(dt) Z
8   bridge maximum M = (x + x' + sqrt((x' - x)^2 - 2 sigma^2 dt log U))
9   / 2
10  barriers in (m, M]: payoff = exp(c) - 1 (hit by the diffusion)
11  m = max(m, M)
12  with prob. lambda/(lambda + q): jump J (Exp(eta_up) up w.p. p,
13   else -Exp(eta_dn)); x = x' + J;
14  barriers in (m, x]: payoff = exp(x) - 1 (passed by the jump)
15  m = max(m, x)
16  otherwise: killed (the time e(q) has come).
17  A path more than 12 below the lowest barrier is stopped.
18 3. Means, standard errors and covariances of the payoffs across paths.
```



References

- [1] Dixit, A. K., & Pindyck, R. S. (1994). *Investment under Uncertainty*. Princeton University Press. <https://doi.org/10.1515/9781400830176>
- [2] McDonald, R., & Siegel, D. (1986). The Value of Waiting to Invest. *Quarterly Journal of Economics*, 101(4), 707–727. <https://doi.org/10.2307/1884175>
- [3] Infrac Foundation (2025). *Price of CO₂ emission allowances EU ETS: primary-market auction results (data: EEX)* [Dataset, open access]. Infrac Foundation, energy.infrac.pl, licence CC BY 4.0.
- [4] Cont, R., & Voltchkova, E. (2005). A finite-difference scheme for option pricing in jump-diffusion and exponential Lévy models. *SIAM Journal on Numerical Analysis*, 43(4), 1596–1626. <https://doi.org/10.1137/s0036142903436186>
- [5] Almendral, A., & Oosterlee, C. W. (2007). On American options under the Variance Gamma process. *Applied Mathematical Finance*, 14(2), 131–152. <https://doi.org/10.1080/13504860600724885>
- [6] Kyprianou, A. E. (2014). *Fluctuations of Lévy Processes with Applications*. Springer. <https://doi.org/10.1007/978-3-642-37632-0>
- [7] Boyarchenko, S. I., & Levendorskiĭ, S. Z. (2002). *Non-Gaussian Merton-Black-Scholes Theory*. World Scientific. <https://doi.org/10.1142/4955>
- [8] Stochastic Analysis Desk (2023). *Resolvent Operators for One-Dimensional Lévy Processes: A Unified Treatment*. Preprint PR-2026-68980035, <https://zamrik.com/research-items/pr-2026-68980035/>, Zamrik Quant Research.
- [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] Alili, L., & Kyprianou, A. E. (2005). Some remarks on first passage of Lévy processes, the American put and pasting principles. *Annals of Applied Probability*, 15(3), 2062–2080. <https://doi.org/10.1214/105051605000000377>
- [11] Crandall, M. G., Ishii, H., & Lions, P.-L. (1992). User’s guide to viscosity solutions of second order partial differential equations. *Bulletin of the American Mathematical Society*, 27(1), 1–67. <https://doi.org/10.1090/s0273-0979-1992-00266-5>
- [12] Jakobsen, E. R., & Karlsen, K. H. (2006). A ”maximum principle for semicontinuous functions” applicable to integro-partial differential equations. *Nonlinear Differential Equations and Applications*, 13, 137–165. <https://doi.org/10.1007/s00030-005-0031-6>
- [13] Barles, G., & Imbert, C. (2008). Second-order elliptic integro-differential equations: Viscosity solutions’ theory revisited. *Annales de l’Institut Henri Poincaré C: Analyse non linéaire*, 25(3), 567–585. <https://doi.org/10.1016/j.anihpc.2007.02.007>



- [14] Novikov, A. A., & Shiryaev, A. N. (2005). On an effective solution of the optimal stopping problem for random walks. *Theory of Probability and Its Applications*, 49(2), 344–354. <https://doi.org/10.1137/S0040585X97981093>
- [15] Kyprianou, A. E., & Surya, B. A. (2005). On the Novikov-Shiryaev optimal stopping problems in continuous time. *Electronic Communications in Probability*, 10, 146–154. <https://doi.org/10.1214/ECP.v10-1144>
- [16] Mordecki, E. (2002). Optimal Stopping and Perpetual Options for Lévy Processes. *Finance and Stochastics*, 6(4), 473–493. <https://doi.org/10.1007/s007800200070>
- [17] 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*, 3(1), 459–478. <https://doi.org/10.1137/110823328>
- [18] Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer. <https://doi.org/10.1007/978-0-387-21617-1>
- [19] Kou, S. G. (2002). A Jump-Diffusion Model for Option Pricing. *Management Science*, 48(8), 1086–1101. <https://doi.org/10.1287/mnsc.48.8.1086.166>
- [20] Kou, S. G., & Wang, H. (2003). First passage times of a jump diffusion process. *Advances in Applied Probability*, 35(2), 504–531. <https://doi.org/10.1239/aap/1051201658>
- [21] European Central Bank (2025). *ECB Data Portal: euro area AAA yield curve, 10-year spot rate* [Dataset, open access]. European Central Bank, data.ecb.europa.eu, licence free reuse with the ECB cited as the source.



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).





Fokker–Planck density of stochastic logistic growth: $\partial_t p = -\partial_x [r x (1 - x/K) p] + \frac{1}{2} \sigma^2 \partial_{xx} (x^2 p)$



ZAMRIK QUANT RESEARCH

The research arm of Zamrik Quant Lab



WP-2026-62242748
zamrik.com