


PREPRINT

Verification Theorems for HJB Equations with α -Stable Lévy Noise

A Viscosity-Solution Framework for Non-Local Generators

Supervised by Dr. T. Zamrik

12-JUL-2024



Phase portrait of the Brusselator: $\dot{x} = a + x^2y - (b+1)x$, $\dot{y} = bx - x^2y$, $a = 0.90$, $b = 3.18$



PR-2026-22911048
zamrik.com



Zamrik Quant Research

The research arm of Zamrik Quant Lab

Copyright

© 2024 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 (2024). ‘Verification Theorems for HJB Equations with α -Stable Lévy Noise: A Viscosity-Solution Framework for Non-Local Generators’. *Zamrik Quant Research Preprint PR-2026-22911048*. Available at <https://zamrik.com/research-items/pr-2026-22911048/>.

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 preprint 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.



PREPRINT

Verification Theorems for HJB Equations with α -Stable Lévy Noise

A Viscosity-Solution Framework for Non-Local Generators

Supervised by Dr. T. Zamrik¹

12-JUL-2024

PR-2026-22911048

Abstract. We establish verification theorems for finite-horizon Hamilton–Jacobi–Bellman equations driven by α -stable Lévy noise, where the classical Itô calculus is replaced by a non-local infinitesimal generator. For $\alpha \in (1, 2)$ and bounded controls, we prove existence and uniqueness of viscosity solutions, derive optimal feedback laws expressible in terms of the resolvent of the non-local operator, and recover the classical Brownian verification theorem as $\alpha \rightarrow 2$. A jump-adapted Monte Carlo scheme for the controlled state is given; on a linear–quadratic subcase with a closed-form value function its statistical error decays as $N^{-1/2}$ at every α (fitted slopes -0.504 to -0.489 for $\alpha \in \{1.3, 1.5, 1.7, 1.9\}$), as plain Monte Carlo must, and its bias is first order in the time step; extrapolated to $\Delta t \rightarrow 0$ it reproduces the closed-form value to 0.11 % (1.0 standard errors).

PR-2026-22911048 · Optimal Control & Decision Desk · Preprint · Published
12-JUL-2024 · Reading time: about 36 min

1. Introduction

The classical Hamilton–Jacobi–Bellman verification programme, developed for controlled diffusion processes driven by Brownian motion [1], [2], relies on Itô’s formula and a second-order infinitesimal generator. When the driving noise is replaced by a pure-jump α -stable Lévy process, the generator becomes non-local and classical $C^{1,2}$ solutions of the resulting partial integro-differential equation (PIDE) may not exist, even in dimension one [3], [4], [5]. The canonical fix — viscosity solutions [6] — extends to the non-local case, but the verification theorem must be re-derived: the martingale argument needs the Itô–Lévy formula in place of Itô’s formula, the polynomial-growth bound must be compatible with the heavy tails of the α -stable law, and the candidate feedback law must be expressible without a classical Hessian.

This paper carries out the re-derivation in a self-contained way and quantifies the result against a tractable benchmark.

We make five contributions.

1. **Comparison principle and uniqueness.** For $\alpha \in (1, 2)$, bounded continuous viscosity sub- and supersolutions of the reduced HJB PIDE satisfy a comparison prin-

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



ciple on $[0, T] \times \mathbb{R}$ under mild regularity of the cost and drift coefficients. Uniqueness of viscosity solutions follows.

2. **Verification theorem.** A smooth viscosity solution of polynomial growth coincides with the value function, and the candidate feedback law obtained by pointwise minimisation of the Hamiltonian is optimal in the admissible class.
3. **Resolvent representation of the optimal feedback.** The optimal feedback is expressible through the resolvent $\mathcal{R}_q = (q\mathbf{I} - \mathcal{L})^{-1}$ of the non-local generator $\mathcal{L}$, evaluated on a single auxiliary function. This places the optimal control in a Banach-algebra setting analogous to the Wiener–Hopf factorisation that drives perpetual-option pricing under jump-diffusion [7], [8].
4. **Numerical scheme.** A jump-adapted Monte Carlo scheme based on the small-jump Gaussian approximation of Asmussen–Rosiński [9] simulates the controlled state under the optimal feedback. Its error is a statistical part of order $N^{-1/2}$ plus a time-step bias of order Δt ; both are measured against a closed form available in a linear–quadratic truncated- α -stable subcase. The index α enters the constant of the statistical error, not its exponent.
5. **Brownian recovery.** As $\alpha \rightarrow 2^-$, the non-local generator converges to the classical second-order operator $\frac{1}{2}\sigma_2^2\sigma^2\partial_x^2$ plus the drift, where $\sigma_2^2 = \lim_{\alpha \rightarrow 2^-} \sigma_R^2 = \frac{1}{2}$ in the normalisation of Section 2, and the verification theorem of this paper reduces to the Fleming–Soner statement.

The paper is organised as follows. Section 2 fixes the control problem. Section 3 derives the non-local HJB PIDE. Section 4 introduces viscosity solutions and states the comparison principle, whose proof is in Appendix A. Section 5 states and proves the verification theorem. Section 6 derives the resolvent representation of the optimal feedback. Section 7 describes the jump-adapted Monte Carlo scheme and states its convergence, whose proof is in Appendix B. Section 8 establishes the Brownian recovery as $\alpha \rightarrow 2^-$. Section 9 reports numerical results on the LQ truncated- α -stable subcase. Sections 10 and 11 discuss limitations, extensions, and conclusions. Appendix C collects the full pseudocode for the three algorithms used in the paper.

2. The control problem under α -stable noise

This section fixes notation and states the control problem solved in the sequel. The state is a one-dimensional controlled SDE driven by a symmetric α -stable Lévy process; the cost is the expected discounted sum of quadratic running and terminal costs.

2.1 State and control

Fix a complete filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \in [0, T]}, \mathbb{P})$ carrying a symmetric α -stable Lévy process $\{L_t^\alpha\}$ of index $\alpha \in (1, 2)$ with Lévy measure



$$\nu_\alpha(dz) = c_\alpha |z|^{-1-\alpha} dz, \quad c_\alpha = \frac{\alpha \Gamma(\frac{1+\alpha}{2})}{2\sqrt{\pi} \Gamma(1 - \frac{\alpha}{2})}, \quad (2.1)$$

truncated to the interval $|z| \leq R$ for some fixed truncation radius $R > 0$. The truncation guarantees that the second moment $\sigma_R^2 := \int_{|z| \leq R} z^2 \nu_\alpha(dz) = 2c_\alpha R^{2-\alpha}/(2-\alpha) < \infty$, which is needed for the LQ cost specification below. *Convention.* With this constant the untruncated measure has characteristic exponent $\int (1 - \cos \xi z) \nu_\alpha(dz) = (|\xi|/2)^\alpha$, which is $2^{-\alpha}$ times the other common normalisation $|\xi|^\alpha$; quadrature of the integral reproduces $2^{-\alpha}$ to 3×10^{-12} for $\alpha \in \{1.3, 1.5, 1.7, 1.9\}$. The scale σ enters once, in the jump σz of the state. The state $X_t \in \mathbb{R}$ evolves as

$$dX_t = (b u_t - \theta X_t) dt + \sigma dL_t^\alpha, \quad X_0 \in \mathbb{R}, \quad (2.2)$$

where $b \in \mathbb{R}$ is the control gain in the drift, $\theta > 0$ is the mean-reversion rate, and $u_t \in U \subseteq \mathbb{R}$ is the control process.

Definition 1 (Admissible control). The control process $u = \{u_t\}_{t \in [0, T]}$ is *admissible* if it is $\{\mathcal{F}_t\}$ -progressively measurable, takes values in a fixed bounded set $U \subset \mathbb{R}$, and satisfies $\mathbb{E} \int_0^T u_t^2 dt < \infty$. The set of admissible controls is denoted $\mathcal{A}$.

2.2 Cost functional

Definition 2 (Cost functional). For an admissible control $u \in \mathcal{A}$ and initial state $X_0 = x$, the *cost functional* is

$$J(t, x; u) = \mathbb{E} \left[\int_t^T e^{-\rho(s-t)} \left(\frac{1}{2} \alpha_c u_s^2 + \frac{1}{2} \beta X_s^2 \right) ds + e^{-\rho(T-t)} \frac{1}{2} \gamma_T X_T^2 \mid X_t = x \right],$$

where $\alpha_c, \beta, \gamma_T > 0$ and $\rho > 0$ is the discount rate.

Assumption 1 (Parameter regularity). $\alpha \in (1, 2)$; $\sigma, \theta, \rho, \alpha_c, \beta, \gamma_T > 0$; $b \in \mathbb{R}$; $R \in (0, \infty)$; $T \in (0, \infty)$; $U \subset \mathbb{R}$ is bounded.

2.3 The value function

The planner solves

$$V(t, x) = \inf_{u \in \mathcal{A}} J(t, x; u), \quad (t, x) \in [0, T] \times \mathbb{R}. \quad (2.3)$$

The first task is to establish that the infimum is finite and attained, and that the value function is continuous. This is standard for the truncated- α -stable case in which the state has finite second moment uniformly in admissible controls.

Proposition 1 (Well-posedness). Under Assumption 1, for every $u \in \mathcal{A}$ and $(t, x) \in [0, T] \times \mathbb{R}$, the SDE for $\{X_s\}_{s \in [t, T]}$ admits a unique càdlàg strong solution with $\mathbb{E} \sup_{s \in [t, T]} X_s^2 < \infty$, the cost functional $J(t, x; u)$ is finite, the value function $V(t, x)$ is finite, and V is continuous on $[0, T] \times \mathbb{R}$.



Proof. The drift is affine in x and bounded in u ; the diffusion coefficient is constant and the truncated Lévy measure has finite second moment, hence the SDE satisfies the standard Lipschitz-and-linear-growth criteria for jump-diffusions [10], [11], Chapter 6, yielding existence, uniqueness, and the moment bound. Finiteness of J follows because the running cost is dominated by $\frac{1}{2} \max(\alpha_c \cdot \sup U^2, \beta)(1 + X_s^2)$ which is integrable. Finiteness of V follows from $V \leq J(t, x; 0)$ for the zero-control. Continuity is proved by the standard dynamic-programming argument [3], [1], using continuity of the running and terminal costs and the moment bound on the state. $\square$

Remark 1 (Why truncation). For pure α -stable noise with $\alpha < 2$, the unbounded Lévy measure forces $\mathbb{E}L_1^2 = \infty$ and the quadratic cost integral diverges. Truncation at $|z| \leq R$ is the simplest fix; equivalent regularisations (tempered stable, CGMY) preserve the heavy-tail structure on $|z| \leq R$ while restoring finite second moment. The viscosity-solution programme below does not require truncation — it is needed only to make the LQ subcase used for numerical validation in Section 9 well-posed.

3. The non-local HJB equation

We derive the Hamilton–Jacobi–Bellman partial integro-differential equation satisfied by the value function and write down its reduced form after pointwise minimisation over the control.

3.1 Dynamic programming and the non-local generator

For a smooth function $\varphi \in C_b^{1,2}([0, T] \times \mathbb{R})$, the infinitesimal generator of the controlled state under a constant control u is the non-local operator

$$\mathcal{L}^u \varphi(t, x) = (bu - \theta x) \partial_x \varphi(t, x) + \int_{\mathbb{R} \setminus \{0\}} [\varphi(t, x + \sigma z) - \varphi(t, x) - \sigma z \partial_x \varphi(t, x) \mathbf{1}_{|z| \leq 1}] \nu_\alpha(dz), \quad (3.1)$$

where the integral is understood as a principal value at $z = 0$ and is absolutely convergent once the second-derivative term is added. The indicator $\mathbf{1}_{|z| \leq 1}$ is the small-jump compensator: under Assumption 1 it could equivalently be set to 1 since the integral is already absolutely convergent on the truncated measure.

Proposition 2 (Dynamic programming principle and HJB). Under Assumption 1, the value function V defined in Section 2 satisfies the dynamic programming principle: for all $(t, x) \in [0, T] \times \mathbb{R}$ and all stopping times $\theta \in [t, T]$,

$$V(t, x) = \inf_{u \in \mathcal{A}} \mathbb{E} \left[\int_t^\theta e^{-\rho(s-t)} L(u_s, X_s) ds + e^{-\rho(\theta-t)} V(\theta, X_\theta) \mid X_t = x \right],$$

where $L(u, x) = \frac{1}{2} \alpha_c u^2 + \frac{1}{2} \beta x^2$. Equivalently, if $V \in C_b^{1,2}$, then V solves the HJB PIDE

$$-\partial_t V(t, x) + \rho V(t, x) = \inf_{u \in U} \{L(u, x) + \mathcal{L}^u V(t, x)\}, \quad V(T, x) = \frac{1}{2} \gamma_T x^2.$$



Proof. The DPP follows from the strong Markov property of the controlled state and standard measurable-selection arguments; see [1], Chapter III, or [3] for the jump-diffusion case. The PIDE is obtained by Itô–Lévy expansion of V under a constant control on a vanishing time interval and taking the limit; see [3], Theorem 3.1. $\square$

3.2 Pointwise minimisation and the reduced HJB PIDE

The Hamiltonian inside the bracket is convex in u and quadratic, so the infimum is attained pointwise. Differentiating with respect to u ,

$$\alpha_c u + b \partial_x V = 0 \implies u^*(t, x) = -\frac{b}{\alpha_c} \partial_x V(t, x), \quad (3.2)$$

provided u^* lies in the bounded admissible set U (we work in the interior throughout). Substituting u^* back gives the *reduced* HJB PIDE.

Theorem 1 (Reduced non-local HJB). Under Assumption 1 and the interior assumption $u^* \in U$, the value function satisfies

$$-\partial_t V + \rho V = -\frac{b^2}{2\alpha_c} (\partial_x V)^2 - \theta x \partial_x V + \frac{1}{2} \beta x^2 + \mathcal{V},$$

where the non-local operator $\mathcal{V}$ acts as

$$\mathcal{V}(t, x) = \int_{\mathbb{R} \setminus \{0\}} [V(t, x + \sigma z) - V(t, x) - \sigma z \partial_x V(t, x) \mathbf{1}_{|z| \leq 1}] \nu_\alpha(dz),$$

with terminal condition $V(T, x) = \frac{1}{2} \gamma_T x^2$.

Proof. Direct substitution of $u^* = -(b/\alpha_c) \partial_x V$ into the bracket on the right-hand side of Proposition 2 and simplification. The nonlinearity in $\partial_x V$ is quadratic; the integral term carries the non-local part of the generator (the drift contribution of the Lévy noise, with the small-jump compensator). $\square$

3.3 Lévy–Khintchine decomposition and a tractable subcase

Remark 2 (LQ subcase). When the noise is symmetric truncated- α -stable as in Section 2 and we restrict to the quadratic ansatz $V(t, x) = \frac{1}{2} A(t) x^2 + C(t)$, the integral $\mathcal{V}$ collapses to the constant $\frac{1}{2} A(t) \sigma^2 \sigma_R^2$ because $\int z \nu_\alpha(dz) = 0$ by symmetry and $\int z^2 \nu_\alpha(dz) = \sigma_R^2$ by truncation. The non-local PIDE reduces to a scalar Riccati ODE for $A(t)$ plus a linear ODE for $C(t)$ — the closed form used in Section 9 to validate the numerical scheme.

Remark 3 (Why the PIDE is non-local). For $V \in C_b^{1,2}$ but not in $C^{1,4}$, the non-local term $\mathcal{V}$ cannot be replaced by a local differential operator. The classical Itô calculus picks up only the second-order derivative through the Brownian quadratic variation; the Lévy noise samples V at all increments $V(x + \sigma z)$ for $z \in (-R, R) \setminus \{0\}$. This is the essential mechanism that forces viscosity solutions in Section 4.



4. Viscosity solutions

Classical $C^{1,2}$ solutions of the reduced HJB PIDE of Theorem 1 are not guaranteed to exist for general data: the non-local term $\mathcal{J}$ imposes additional regularity that is hard to verify a priori [3], [4]. The viscosity-solution framework [6] extends to the non-local case provided the test functions are equipped with the non-local correction [12].

4.1 Definition

Definition 3 (Viscosity sub- and supersolution). A locally bounded upper semicontinuous function $V: [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$ is a *viscosity subsolution* of the reduced HJB PIDE of Theorem 1 if, for every $\varphi \in C_b^{1,2}$ and every local maximum $(t_0, x_0) \in (0, T) \times \mathbb{R}$ of $V - \varphi$, one has

$$-\partial_t \varphi(t_0, x_0) + \rho V(t_0, x_0) + \frac{b^2}{2\alpha_c} (\partial_x \varphi(t_0, x_0))^2 + \theta x_0 \partial_x \varphi(t_0, x_0) - \frac{1}{2} \beta x_0^2 - \mathcal{J}^{\varphi, V}(t_0, x_0) \leq 0,$$

where the non-local term is split as $\mathcal{J}^{\varphi, V} = \mathcal{J}^{1, \varphi} + \mathcal{J}^{2, V}$ with

$$\begin{aligned} \mathcal{J}^{1, \varphi}(t_0, x_0) &= \int_{|z| \leq 1} [\varphi(t_0, x_0 + \sigma z) - \varphi(t_0, x_0) - \sigma z \partial_x \varphi(t_0, x_0)] \nu_\alpha(dz), \\ \mathcal{J}^{2, V}(t_0, x_0) &= \int_{1 < |z| \leq R/\sigma} [V(t_0, x_0 + \sigma z) - V(t_0, x_0)] \nu_\alpha(dz). \end{aligned}$$

A locally bounded lower semicontinuous function V is a *viscosity supersolution* if the analogous inequality with $\leq$ replaced by $\geq$ holds at every local minimum. A *viscosity solution* is both.

The split $\mathcal{J} = \mathcal{J}^{1, \varphi} + \mathcal{J}^{2, V}$ is the standard Barles–Imbert decomposition [4]: small jumps use the test function (where regularity is controlled), large jumps use the candidate V itself (where only semicontinuity is available). The split is necessary because the small-jump compensator $-\sigma z \partial_x \varphi \cdot \mathbf{1}_{|z| \leq 1}$ requires $\partial_x \varphi$ but the large-jump increments do not.

4.2 Comparison principle

Theorem 2 (Comparison). Under Assumption 1, let $\underline{V}$ be a viscosity subsolution and $\overline{V}$ a viscosity supersolution of the reduced HJB PIDE of Theorem 1 on $[0, T] \times \mathbb{R}$, both with polynomial growth in x . If $\underline{V}(T, x) \leq \overline{V}(T, x)$ for all $x \in \mathbb{R}$, then $\underline{V}(t, x) \leq \overline{V}(t, x)$ for all $(t, x) \in [0, T] \times \mathbb{R}$.

Proof. Doubling-the-variables argument with the non-local correction; see Appendix A for the full proof. The key step is the Barles–Imbert maximum principle for semicontinuous functions [12], [4], applied to the auxiliary function $(\underline{V}(t, x) - \overline{V}(s, y)) - \delta(t + s) - (|x - y|^2 + \varepsilon|x|^2 + \varepsilon|y|^2)/(2\eta)$ on $[0, T]^2 \times \mathbb{R}^2$. The penalisation in ε handles unboundedness of the domain; the penalisation in η localises the contact point. The non-local integral terms are estimated by splitting at the test-function radius and using polynomial growth on the large-jump side. $\square$



4.3 Existence and uniqueness

Theorem 3 (Existence). Under Assumption 1, the value function V of Section 2 is a viscosity solution of the reduced HJB PIDE of Theorem 1.

Proof. Standard application of the dynamic programming principle (Proposition 2) combined with the Itô–Lévy formula on smooth test functions, following the recipe of [3], Theorem 4.1, adapted to the truncated- α -stable case by the change of variables $z \mapsto \sigma z$. $\square$

Corollary 1 (Uniqueness). Under Assumption 1, the value function V is the unique viscosity solution of the reduced HJB PIDE within the class of continuous functions of polynomial growth, with terminal condition $V(T, x) = \frac{1}{2}\gamma_T x^2$.

Proof. Apply Theorem 2 to any two viscosity solutions V_1, V_2 of polynomial growth with the same terminal data: $V_1 \leq V_2$ and $V_2 \leq V_1$, hence $V_1 \equiv V_2$. Theorem 3 then identifies the unique viscosity solution with the value function. $\square$

5. Verification theorem

The viscosity-solution characterisation of Section 4 identifies the value function uniquely but does not by itself yield a useful feedback law: the candidate $u^*(t, x) = -(b/\alpha_c)\partial_x V(t, x)$ requires $\partial_x V$, which is not part of the viscosity-solution data. Verification closes the gap: if a smooth function happens to be a viscosity solution of polynomial growth, then it is the value function and the candidate feedback is optimal.

5.1 Statement

Theorem 4 (Verification). Let $\widehat{V} \in C^{1,2}([0, T] \times \mathbb{R}) \cap C([0, T] \times \mathbb{R})$ be a viscosity solution of the reduced HJB PIDE of Theorem 1 with terminal condition $\widehat{V}(T, x) = \frac{1}{2}\gamma_T x^2$. Define the candidate feedback law

$$\widehat{u}(t, x) := -\frac{b}{\alpha_c} \partial_x \widehat{V}(t, x),$$

and assume the closed-loop SDE

$$d\widehat{X}_t = (b\widehat{u}(t, \widehat{X}_t) - \theta\widehat{X}_t)dt + \sigma dL_t^\alpha, \quad \widehat{X}_0 = x,$$

has a unique strong càdlàg solution that satisfies $\mathbb{E} \sup_{t \in [0, T]} \widehat{X}_t^2 < \infty$, and $\widehat{u}(t, \widehat{X}_t) \in U$ almost surely. Assume further the polynomial-growth bound

$$|\widehat{V}(t, x)| + |\partial_x \widehat{V}(t, x)| \leq K(1 + x^2) \quad \text{for some } K > 0, \text{ uniformly in } t \in [0, T].$$

Then $\widehat{V} = V$ on $[0, T] \times \mathbb{R}$ and $\widehat{u}$ is an optimal feedback control.



5.2 Proof

Proof. Fix $(t, x) \in [0, T] \times \mathbb{R}$ and let $u \in \mathcal{A}$ be an arbitrary admissible control generating the state $\{X_s\}_{s \in [t, T]}$ with $X_t = x$. The discounted candidate $e^{-\rho(s-t)} \widehat{V}(s, X_s)$ is a semimartingale; the Itô–Lévy formula gives

$$e^{-\rho(T-t)} \widehat{V}(T, X_T) - \widehat{V}(t, x) = \int_t^T e^{-\rho(s-t)} \mathcal{H}^{u_s} \widehat{V}(s, X_{s-}) ds + M_T,$$

with the discounted infinitesimal operator

$$\mathcal{H}^u \widehat{V} := -\rho \widehat{V} + \partial_t \widehat{V} + (bu - \theta x) \partial_x \widehat{V} + \mathcal{J} \widehat{V},$$

and $\{M_s\}$ the stochastic integral against the compensated Poisson random measure of L^α ; the polynomial-growth bound and the second-moment bound on X (which follows from boundedness of u) ensure that $\{M_s\}$ is a square-integrable martingale by Burkholder–Davis–Gundy applied to the Poisson compensator. Adding the running cost on both sides,

$$\widehat{V}(t, x) = \mathbb{E} \left[\int_t^T e^{-\rho(s-t)} L(u_s, X_s) ds + e^{-\rho(T-t)} \Psi(X_T) \right] - \mathbb{E} \int_t^T e^{-\rho(s-t)} [\mathcal{H}^{u_s} \widehat{V} + L] ds,$$

where $L(u, x) = \frac{1}{2} \alpha_c u^2 + \frac{1}{2} \beta x^2$, $\Psi(x) = \frac{1}{2} \gamma_T x^2$, and the martingale expectation has vanished. The reduced HJB PIDE gives $\mathcal{H}^u \widehat{V} + L \geq 0$ pointwise (by the infimum definition), with equality if and only if $u = \hat{u}$. Therefore

$$\widehat{V}(t, x) \leq J(t, x; u) \quad \forall u \in \mathcal{A},$$

with equality at $u = \hat{u}$. Taking the infimum on the right gives $\widehat{V}(t, x) = V(t, x)$; the candidate $\hat{u}$ attains the infimum and is therefore optimal. $\square$

Corollary 2 (Optimal feedback). The optimal control problem of Section 2 has an optimal feedback law in the bounded admissible set U , given by

$$u^*(t, x) = -\frac{b}{\alpha_c} \partial_x V(t, x),$$

whenever the right-hand side lies in U ; otherwise the optimal control is the boundary value of U at which $L(u, x) + bu \partial_x V(t, x)$ is minimised.

Remark 4 (Where the polynomial-growth bound enters). The bound serves three purposes. First, it ensures $\widehat{V}(T, X_T) \in L^1$ so the terminal term in the Itô–Lévy expansion can be evaluated under expectation. Second, it ensures $\partial_x \widehat{V} \in L^2$ along the state, which is needed for the martingale property of $\{M_s\}$ via BDG. Third, it enforces uniqueness through Theorem 2 — without polynomial growth, comparison can fail because the test-function class is too wide.

Remark 5 (Brownian limit). In the limit $\alpha \rightarrow 2^-$, the non-local term $\mathcal{J} \widehat{V}$ reduces to $\frac{1}{2} \sigma_2^2 \sigma^2 \partial_x^2 \widehat{V}$ with $\sigma_2^2 = \frac{1}{2}$ (Section 8), the Itô–Lévy formula reduces to Itô’s formula, and the verification theorem above reduces to the classical Fleming–Soner statement [1], Chapter IV.



6. Optimal feedback via the resolvent

Corollary 2 expresses the optimal feedback as $u^*(t, x) = -(b/\alpha_c) \partial_x V(t, x)$. The value function itself solves the non-local HJB PIDE of Theorem 1, so computing u^* reduces to computing $\partial_x V$. In the linear–quadratic subcase the ansatz $V(t, x) = \frac{1}{2}A(t)x^2 + C(t)$ gives a closed form; outside the LQ class, the natural object is the *resolvent* of the non-local generator.

6.1 The resolvent operator

Definition 4 (Resolvent). For $q > 0$, the resolvent $\mathcal{R}_q$ of the non-local generator $\mathcal{L}$ is the operator

$$\mathcal{R}_q = (qI - \mathcal{L})^{-1},$$

acting on the Banach space $C_0(\mathbb{R})$ of continuous functions vanishing at infinity, with

$$(\mathcal{R}_q f)(x) = \int_0^\infty e^{-qt} (\mathbb{E}^x f(X_t)) dt,$$

where $\mathbb{E}^x$ denotes expectation under the law of $\{X_t\}$ started at x . The resolvent is bounded with $\|\mathcal{R}_q\| \leq 1/q$ and satisfies the Hille–Yosida identity $\mathcal{R}_q - \mathcal{R}_p = (p - q) \mathcal{R}_q \mathcal{R}_p$.

Without drift ($\theta = 0, u \equiv 0$) the generator is translation invariant and the resolvent is the Fourier multiplier

$$\widehat{\mathcal{R}_q}(\xi) = \frac{1}{q + \psi_R(\sigma\xi)}, \quad \psi_R(\xi) = \int_{|z| \leq R} (1 - \cos \xi z) \nu_\alpha(dz), \quad (6.1)$$

with $\psi_R(\xi) \rightarrow (|\xi|/2)^\alpha$ as $R \rightarrow \infty$ in the convention of Section 2. A mean-reverting or state-dependent drift breaks translation invariance, and the resolvent is then no longer a multiplier; for a linear closed loop it acts on linear functions explicitly, which is what Theorem 5 uses.

6.2 Representation of the optimal feedback

Theorem 5 (Resolvent representation). Suppose Theorem 4 holds, that the optimal feedback u^* is C^1 in x with $\partial_x u^*$ bounded, and let $\{X_s\}_{s \in [t, T]}$ be the closed-loop state started at $X_t = x$. Write $h(s, x) = \frac{1}{2}\alpha_c (u^*(s, x))^2 + \frac{1}{2}\beta x^2$ for the realised cost-rate, $\Psi(x) = \frac{1}{2}\gamma_T x^2$, and

$$\Phi_{t,s} = \exp\left(\int_t^s (b \partial_x u^*(r, X_r) - \theta) dr\right)$$

for the derivative of the closed-loop flow with respect to its starting point. Then

$$\partial_x V(t, x) = \mathbb{E}^{t,x} \left[\int_t^T e^{-\rho(s-t)} \Phi_{t,s} \partial_x h(s, X_s) ds + e^{-\rho(T-t)} \Phi_{t,T} \Psi'(X_T) \right].$$

When the closed loop is linear and time-homogeneous, $u^*(x) = -kx$ with $\kappa := \theta + bk > 0$, then $\Phi_{t,s} = e^{-\kappa(s-t)}$ and on the infinite horizon the representation is the resolvent of the



closed-loop generator at the shifted rate $\rho + \kappa$:

$$\partial_x V(x) = (\mathcal{R}_{\rho+\kappa}^{u^*} \partial_x h)(x).$$

Proof. By Theorem 4, $V(t, x) = J(t, x; u^*)$, the expected discounted cost along the closed loop. The noise is additive, so the closed-loop flow is differentiable in its starting point with derivative $\partial_x X_s = \Phi_{t,s}$, which solves $d\Phi_{t,s} = (b \partial_x u^*(s, X_s) - \theta) \Phi_{t,s} ds$, $\Phi_{t,t} = 1$; it is bounded by $e^{(b\|\partial_x u^*\|_\infty + \theta)(T-t)}$. Differentiating under the expectation (dominated convergence, with the second-moment bound of Proposition 1 and the polynomial growth of $\partial_x h$ and Ψ') and the chain rule give the first formula. In the time-homogeneous linear case $\Phi_{t,s} = e^{-\kappa(s-t)}$ is deterministic, and letting $T \rightarrow \infty$ the first formula becomes $\int_0^\infty e^{-(\rho+\kappa)s} \mathbb{E}^x[\partial_x h(X_s)] ds$, which is Definition 4 with $q = \rho + \kappa$. $\square$

Remark 6 (Wiener–Hopf parallel). For perpetual stopping problems on the same state $\{X_t\}$, the resolvent admits a Wiener–Hopf factorisation $\mathcal{R}_q = \mathcal{R}_q^+ \mathcal{R}_q^-$ into one-sided resolvents; this is the engine of pure-jump option pricing [7], [8]. The verification theorem above does not require the factorisation directly, but the resolvent representation places the optimal-control problem in the same operator-algebraic setting and opens the door to closed-form computations whenever the symbol of $\mathcal{R}_q$ is rational in $|\xi|^\alpha$ (the CGMY and tempered-stable cases).

Remark 7 (Algorithmic content). The representation gives $\partial_x V$, and hence the feedback, at a single state without solving the PIDE on a grid: simulate the closed loop, carry $\Phi_{t,s}$ along each path, and average the bracket — the pathwise-derivative estimator that Algorithm 5 accumulates next to the cost. In the LQ subcase both the quadrature of the representation and the pathwise estimator are compared with the closed-form gain in Section 9.

7. Numerical scheme — jump-adapted Monte Carlo

For numerical work on the controlled state and the value function, we use the *jump-adapted Monte Carlo* scheme of Asmussen–Rosiński [9]. The scheme splits the truncated- α -stable Lévy measure into a small-jump part (handled by a Gaussian approximation) and a large-jump part (handled by a Poisson clock with exact jump-size sampling).

7.1 Scheme overview

Fix a small-jump truncation level $\varepsilon \in (0, R)$. The Lévy measure splits as $\nu_\alpha = \nu_\alpha^{\leq \varepsilon} + \nu_\alpha^{> \varepsilon}$ where the small part is supported on $|z| \leq \varepsilon$ and the large part on $\varepsilon < |z| \leq R$. The large-jump intensity is

$$\lambda_\varepsilon = \int_{\varepsilon < |z| \leq R} \nu_\alpha(dz) = \frac{2c_\alpha}{\alpha} (\varepsilon^{-\alpha} - R^{-\alpha}). \quad (7.1)$$

Large jumps are sampled from the truncated stable density restricted to $\varepsilon \leq |z| \leq R$ by inverse-CDF. Small jumps are approximated by a Gaussian increment with mean zero and variance



$$\sigma_\varepsilon^2 = \int_{|z| \leq \varepsilon} z^2 \nu_\alpha(dz) = \frac{2c_\alpha}{2-\alpha} \varepsilon^{2-\alpha}. \quad (7.2)$$

Between large-jump times, the drift is integrated by explicit Euler.

Algorithm 1 — Jump-adapted single-step transition

Compute the next state at the next jump time or at the maximum sub-step, whichever comes first. Inputs: parameters $(b, \theta, \sigma, \alpha, R)$; truncation ε ; maximum sub-step Δt ; current state x ; current time t ; control feedback u^* ; random number generator. Output: post-step state, post-step time. *Full pseudocode in Appendix C, Algorithm 4.*

Algorithm 2 — Full trajectory simulator under optimal feedback

Inputs: parameters; ε , Δt ; horizon T ; initial state X_0 ; path count N ; seed. Output: N paths $\{\widehat{X}_t^{(m)}\}_{t \in [0, T]}$ and the accumulated discounted cost $\widehat{J}_m$ along each path. *Full pseudocode in Appendix C, Algorithm 5.*

7.2 Convergence

Proposition 3 (Convergence of the empirical value). Let $\widehat{V}_{N, \varepsilon, \Delta t}(0, x) = \frac{1}{N} \sum_{m=1}^N \widehat{J}_m$ be the empirical value produced by Algorithm 2 with N paths, small-jump truncation ε , and outer step Δt . Under Assumption 1 and the further mild regularity $\partial_x V \in L^\infty([0, T] \times \mathbb{R})$, there exists a constant $K > 0$ independent of $(N, \varepsilon, \Delta t)$ such that

$$\mathbb{E}|\widehat{V}_{N, \varepsilon, \Delta t}(0, x) - V(0, x)| \leq \frac{\sigma_J}{\sqrt{N}} + K(\Delta t + \varepsilon^{3-\alpha}),$$

where σ_J^2 is the variance of one path's cost. The first term is statistical; the other two are biases that do not shrink with N . In the LQ subcase the ε -term vanishes: the expected cost depends on the noise only through its mean and variance, which the Gaussian small-jump increment matches exactly.

Proof. The error decomposes additively into three contributions: (i) Monte Carlo, $\sigma_J N^{-1/2}$ by the Cauchy–Schwarz inequality applied to the mean of N i.i.d. costs $\{\widehat{J}_m\}$ (finite variance because the truncated state has finite fourth moment); (ii) Euler discretisation and left-point accumulation of the cost, $O(\Delta t)$ by the Lipschitz-and-linear-growth jump-diffusion estimates of [13], Chapter 6; (iii) the small-jump Gaussian approximation [9]: the Gaussian increment matches the first two moments of the small-jump increment, so the weak error is governed by the third absolute moment $\int_{|z| \leq \varepsilon} |z|^3 \nu_\alpha(dz) = 2c_\alpha \varepsilon^{3-\alpha}/(3-\alpha)$. Full details in Appendix B. $\square$

Remark 8 (What α changes). Plain Monte Carlo over i.i.d. paths with a finite-variance cost cannot converge faster than $N^{-1/2}$, and does not converge slower: the exponent is $1/2$ for every α . The index enters only through constants — the path-cost spread σ_J , and the work per path, which grows with the large-jump intensity λ_ε . Section 9 measures the slope at $\alpha \in \{1.3, 1.5, 1.7, 1.9\}$ and the bias separately.



8. Brownian recovery ($\alpha \rightarrow 2^-$)

The α -stable framework reduces to the classical Brownian case in the limit $\alpha \rightarrow 2^-$. The non-local generator converges to the second-order local operator, viscosity solutions converge to classical solutions, and the verification theorem of Section 5 reduces to the Fleming–Soner statement [1]. We make this precise.

8.1 Generator convergence

For each $\alpha \in (1, 2)$ write $\mathcal{L}_\alpha$ for the truncated- α -stable generator of Section 3 and let $\mathcal{L}_2 = \frac{1}{2}\sigma_2^2\sigma^2\partial_x^2 + (bu - \theta x)\partial_x$, $\sigma_2^2 := \lim_{\alpha \rightarrow 2^-} \sigma_R^2 = \frac{1}{2}$, be the classical Brownian generator with the same drift. (Near $\alpha = 2$, $\Gamma(1 - \alpha/2) \sim 2/(2 - \alpha)$, so $c_\alpha \sim (2 - \alpha)/4$ and $\sigma_R^2 = 2c_\alpha R^{2-\alpha}/(2 - \alpha) \rightarrow \frac{1}{2}$ for every R ; with the normalisation $|\xi|^\alpha$ the limit would be 2.)

Proposition 4 (Pointwise generator convergence). For every $\varphi \in C_b^2(\mathbb{R})$ and every $x \in \mathbb{R}$,

$$\mathcal{L}_\alpha \varphi(x) \rightarrow \mathcal{L}_2 \varphi(x) \quad \text{as } \alpha \rightarrow 2^-,$$

with the rate $O(2 - \alpha)$ in the supremum norm over compact sets.

Proof. Write the non-local part of $\mathcal{L}_\alpha$ as

$$\mathcal{J}_\alpha \varphi(x) = \int_{|z| \leq R} [\varphi(x + \sigma z) - \varphi(x) - \sigma z \varphi'(x) \mathbf{1}_{|z| \leq 1}] c_\alpha |z|^{-1-\alpha} dz.$$

A Taylor expansion of $\varphi(x + \sigma z)$ to second order yields, after using the symmetry of the Lévy measure (so the first-order term vanishes),

$$\mathcal{J}_\alpha \varphi(x) = \frac{1}{2}\sigma^2 \varphi''(x) \int_{|z| \leq R} |z|^2 c_\alpha |z|^{-1-\alpha} dz + (\text{third-order remainder}).$$

The integral is $\sigma_R^2 = 2c_\alpha R^{2-\alpha}/(2 - \alpha)$, and $\sigma_R^2 - \sigma_2^2 = O(2 - \alpha)$ by the expansion of c_α and of $R^{2-\alpha}$ at $\alpha = 2$ (see [10], Chapter 3, for the stable constants). The third-order remainder is bounded by $\frac{1}{6}\sigma^3 \|\varphi'''\|_\infty \int_{|z| \leq R} |z|^3 \nu_\alpha(dz) = \frac{1}{3}\sigma^3 \|\varphi'''\|_\infty c_\alpha R^{3-\alpha}/(3 - \alpha) = O(2 - \alpha)$. $\square$

8.2 Viscosity-solution convergence

Theorem 6 (Brownian recovery of viscosity solutions). Let V_α be the unique viscosity solution of the reduced HJB PIDE of Theorem 1 under the truncated- α -stable noise with parameters as in Section 2. Let V_2 be the unique viscosity (in this case classical $C^{1,2}$) solution of the corresponding Brownian HJB equation

$$-\partial_t V_2 + \rho V_2 = -\frac{b^2}{2\alpha_c} (\partial_x V_2)^2 - \theta x \partial_x V_2 + \frac{1}{2}\beta x^2 + \frac{1}{2}\sigma_2^2 \sigma^2 \partial_x^2 V_2,$$

with the same terminal condition. Then $V_\alpha \rightarrow V_2$ locally uniformly on $[0, T] \times \mathbb{R}$ as $\alpha \rightarrow 2^-$.

Proof. Apply the half-relaxed limits method of [4], [12]: define $\bar{V}(t, x) = \limsup_{\alpha \rightarrow 2^-}^* V_\alpha$



and $\underline{V}(t, x) = \liminf_{*, \alpha \rightarrow 2^-} V_\alpha$. By Proposition 4 and the stability of viscosity solutions under operator limits, $\bar{V}$ is a viscosity subsolution and $\underline{V}$ is a viscosity supersolution of the Brownian HJB. Both have polynomial growth and agree at $t = T$ with the terminal data. The comparison principle (Theorem 2, in the Brownian limit) gives $\bar{V} \leq V_2 \leq \underline{V}$, and since $\underline{V} \leq \bar{V}$ by construction, equality holds. Local uniform convergence follows. $\square$

Remark 9 (Verification recovery). The verification theorem (Theorem 4), under the Brownian limit, reduces to the classical statement of [1], Chapter IV, with the Itô–Lévy formula collapsing to Itô’s formula because the non-local term vanishes when the candidate is sufficiently smooth.

Remark 10 (Numerical perspective). The $\alpha \rightarrow 2^-$ limit is also visible in the Monte Carlo scheme of Section 7: as α approaches 2, the small-jump variance σ_ε^2 dominates the large-jump intensity λ_ε , and the scheme reduces to a pure Gaussian Euler step — which is the standard Euler–Maruyama scheme for Brownian-driven SDEs. The statistical error is $N^{-1/2}$ on both sides of the limit; only its constant moves.

9. Numerical results

We check the theory on the LQ truncated- α -stable subcase of Remark 2, where the value function is known in closed form. Every number below is computed by the scripts that accompany the paper (`hjb_stable.py`, `mc_jump_adapted.py`, `checks.py`), and every closed form is checked by a second, independent method.

9.1 Closed form in the LQ subcase

The quadratic ansatz $V(t, x) = \frac{1}{2}A(t)x^2 + C(t)$ inserted into the reduced HJB PIDE of Theorem 1 gives, after the non-local term collapses to $\frac{1}{2}A(t)\sigma^2\sigma_R^2$, the Riccati system in time-to-horizon $s = T - t$:

$$\frac{dA}{ds} = -(\rho + 2\theta)A - \frac{b^2}{\alpha_c}A^2 + \beta, \quad A(0) = \gamma_T, \quad (9.1)$$

$$\frac{dC}{ds} = -\rho C + \frac{1}{2}\sigma^2\sigma_R^2A, \quad C(0) = 0. \quad (9.2)$$

The Riccati equation has the Möbius closed form with rate $\lambda = \sqrt{(\rho + 2\theta)^2 + 4(b^2/\alpha_c)\beta}$ and steady-state value $A_\infty = ((-\rho - 2\theta) + \lambda)/(2b^2/\alpha_c)$; $C(t)$ follows by adaptive quadrature. As a second method we integrate the system with an adaptive eighth-order Runge–Kutta solver: on a 501-point grid of $[0, T]$ the two agree to 8×10^{-13} for A and 1×10^{-13} for C .

9.2 Calibration

For the headline numbers we adopt $\alpha = 1.7$, $\sigma = 1.0$, $\theta = 0.10$, $b = 1.0$, $\rho = 0.05$, $\alpha_c = \beta = \gamma_T = 1.0$, $R = 5.0$, $T = 5.0$. At these values the truncated second moment is $\sigma_R^2 = 0.7423$ (closed form and quadrature of $\int z^2 \nu_\alpha(dz)$ agree to 4×10^{-15}); the Riccati steady state is $A_\infty = 0.8828$; the decay rate is $\lambda = 2.0156$ with half-life $\ln 2/\lambda = 0.344$. At



time-to-horizon $s = 1.5$ the gap $(A - A_\infty)/(\gamma_T - A_\infty)$ has fallen to 0.046 (it decays like $e^{-\lambda s}$, and $e^{-1.5\lambda} = 0.049$), so the gain is in steady state for $s > 1.5$, that is for $t < T - 1.5$. At the reference state $x_{\text{ref}} = 1$ the closed-form value is $V(0, 1) = 1.9077$ and the optimal control is $u^*(0, 1) = -0.8828$.

9.3 The closed form solves the PIDE

We evaluate the non-local term $\mathcal{N}$ of Theorem 1 by direct quadrature over z , without the collapse of Remark 2, and the time derivative by finite differences of the closed form. The residual of the reduced HJB PIDE of Theorem 1 at 30 points $(t, x) \in \{0, 1, 2.5, 4, 4.9\} \times \{-2, -1, 0, 0.5, 1, 2\}$ is at most 2×10^{-9} , the size of the finite-difference error. The closed form therefore solves the HJB equation of Proposition 2 with the pointwise minimiser u^* , and the simulations below show that it is the value of the control problem.

9.4 Comparison and uniqueness

Theorem 2 says that ordered terminal data give ordered solutions. We solve the LQ problem with $\gamma_T \in \{0.5, 1, 2\}$ (classical, hence viscosity, solutions) on a 51×61 grid of $[0, T] \times [-3, 3]$. The differences $V_{\gamma_T=1} - V_{\gamma_T=0.5}$ and $V_{\gamma_T=2} - V_{\gamma_T=1}$ are nonnegative at every grid point; they vanish only at $(T, 0)$, where the terminal data coincide, and are at least 0.0025 elsewhere, as Theorem 2 and its restatement Theorem 7 require. The value of the control problem, estimated by simulation in Table 1 with no PIDE involved, agrees with the closed-form solution, as Theorem 3 and the uniqueness of Corollary 1 require.

9.5 The verification theorem

Theorem 4 says that $V \leq J(u)$ for every admissible u , with equality at u^* . Table 1 compares the optimal feedback with four others, $u = -m(b/\alpha_c)A(t)x$ for $m \in \{0, 0.5, 1.5, 2\}$ ($m = 0$ is no control). Each cost is computed twice: from the linear ODE that the Feynman–Kac formula gives for a linear closed loop, and by the jump-adapted Monte Carlo of Algorithm 2 with 10^6 paths at $\varepsilon = 0.05$, $\Delta t = 0.0125$. At $m = 1$ the cost ODE returns $V(0, 1)$ to 8×10^{-15} , a third route to the value. Every other feedback costs more, by margins of 81 to 2023 Monte Carlo standard errors. The simulated costs exceed the cost ODE by 0.14 % to 1.21 %: the time-step bias of the scheme, measured below. The LQ feedback is linear in x , so U is unbounded here; the martingale step of the proof needs only the quadratic growth and the second-moment bound of Proposition 1, which the truncated noise supplies: the simulated $\mathbb{E} \sup_s X_s^2$ is 2.750 under u^* and 5.790 without control.



Table 1: Theorem 4 at $(t, x) = (0, 1)$, $\alpha = 1.7$: the cost of five linear feedbacks from the Feynman–Kac cost ODE and from 10^6 simulated paths each ($\varepsilon = 0.05$, $\Delta t = 0.0125$; $\pm$ one standard error). The optimal feedback is the cheapest, and its cost is $V(0, 1)$. The last column is the time-step bias of the scheme, not noise.

Feedback $u = -m(b/\alpha_c)A(t)x$	$J(0, 1; u)$, cost ODE	$J(0, 1; u)$, Monte Carlo	MC – ODE
$m = 0$ (no control)	5.39741	5.41442 ± 0.00669	+0.01701
$m = 0.5$	2.23604	2.23907 ± 0.00222	+0.00302
$m = 1$ (optimal, u^*)	1.90766	1.92203 ± 0.00173	+0.01437
$m = 1.5$	2.04672	2.06252 ± 0.00176	+0.01579
$m = 2$	2.33868	2.36695 ± 0.00195	+0.02827

9.6 Value function and optimal-policy contours

Figure 1 (left) shows the value function $V(t, x)$ over $(t, x) \in [0, T] \times [-3, 3]$ at $\alpha = 1.7$; it is convex in x and varies little in t . On the right, $V(0, x)$ is set against the Monte Carlo value under u^* at five states and against the costs of two non-optimal feedbacks, which lie above it everywhere (Theorem 4).




Figure 1: Value function and verification.

Figure 1. Left: value function $V(t, x) = \frac{1}{2}A(t)x^2 + C(t)$ at $\alpha = 1.7$ with the calibration above; white lines are level curves. Right: $V(0, x)$ in closed form (solid), the Monte Carlo value under u^* at five states (2.5×10^5 paths each, $\Delta t = 0.0125$, 95 % intervals), and the costs of the feedbacks with the gain halved and multiplied by 1.5 (grey, from the cost ODE).



Figure 2 shows the optimal feedback $u^*(t, x) = -(b/\alpha_c) A(t) x$. It is linear in x with a negative slope that drives the state back toward the origin. The gain $A(t)$ sits at its steady-state value $A_\infty = 0.883$ for most of the horizon and rises to $A(T) = \gamma_T = 1$ in the last 1.5 time units, where the terminal cost takes over; it does not relax toward A_∞ as $t \rightarrow T$, it leaves it.

The representation of Theorem 5 is checked on the same gain. In the LQ case $\Phi_{t,s} = \exp(-\int_t^s (\theta + b^2 A(r)/\alpha_c) dr)$ and $\mathbb{E}^{t,x}[\partial_x h(s, X_s)] = (\beta + b^2 A(s)^2/\alpha_c) \Phi_{t,s} x$, so Theorem 5 states $A(t) = \int_t^T e^{-\rho(s-t)} \Phi_{t,s}^2 (\beta + b^2 A(s)^2/\alpha_c) ds + e^{-\rho(T-t)} \Phi_{t,T}^2 \gamma_T$. Evaluated by quadrature at $t \in \{0, 2.5, 4, 4.5, 4.9\}$ it reproduces the closed-form gain to 3×10^{-16} . The pathwise estimator of Remark 7, averaged over the simulated closed loop and extrapolated to $\Delta t \rightarrow 0$ as in the next subsection, gives $\partial_x V(0, 1) = 0.8830 \pm 0.0006$ against $A(0) = 0.8828$. In the stationary limit the closed-loop rate is $\kappa_\infty = \theta + b^2 A_\infty/\alpha_c = 0.9828$, and the resolvent at the shifted rate $\rho + \kappa_\infty$ gives $(\beta + b^2 A_\infty^2/\alpha_c)/(\rho + 2\kappa_\infty) = 0.8828 = A_\infty$; the resolvent at the unshifted rate ρ would give 1.7228.




Figure 2: Optimal feedback and its gain.

Figure 2. Left: optimal feedback $u^*(t, x) = -(b/\alpha_c) A(t) x$ at $\alpha = 1.7$. Right: the gain $A(t)$ in closed form (solid), from an ODE solve of the Riccati equation (crosses), from the representation of Theorem 5 by quadrature (circles), and from the pathwise Monte Carlo estimator extrapolated to $\Delta t \rightarrow 0$ (square, 95 % interval); the dashed line is A_∞ .

9.7 Monte Carlo convergence

We run the scheme as Algorithm 2 prints it: the controlled drift $b u^* - \theta x$, the running cost $\frac{1}{2} \alpha_c u^{*2} + \frac{1}{2} \beta x^2$ accumulated at the left end of each sub-step, and the terminal cost $\frac{1}{2} \gamma_T X_T^2$. With $\varepsilon = \Delta t = 0.05$ and 4×10^6 paths at each $\alpha \in \{1.3, 1.5, 1.7, 1.9\}$, we split the paths into batches of $N \in \{100, 250, 1000, 2500, 10^4\}$ and measure the root-mean-square deviation of the batch means from the mean of all the paths — the statistical error of Proposition 3 (and of its restatement, Theorem 8). Figure 3 (left) shows it on



log-log axes. The least-squares slopes are -0.504 ± 0.007 , -0.489 ± 0.005 , -0.493 ± 0.012 , -0.501 ± 0.006 for $\alpha = 1.3, 1.5, 1.7, 1.9$ (standard errors from eight independent sub-samples of the paths), against $-1/2$. At every N the root-mean-square error is within 5 % of the central-limit value $\sigma_J/\sqrt{N}$, so the four curves differ only by the constant σ_J ($\sigma_J = 2.959, 2.345, 1.744, 1.117$).

The bias does not shrink with N . At $\varepsilon = \Delta t = 0.05$ the mean of the 4×10^6 paths exceeds the closed form by 0.0618, 0.0504, 0.0415, 0.0367 for $\alpha = 1.3, 1.5, 1.7, 1.9$ (standard errors 0.0015 or less). It is the time-step bias of the Euler drift and the left-point cost (Lemma 7). Figure 3 (right) shows it against Δt at $\alpha = 1.7$ for two truncation levels, 10^6 paths per point; on log-log axes it falls with $\Delta t \leq 0.05$ at slopes 0.91 ± 0.07 ($\varepsilon = 0.05$) and 1.12 ± 0.08 ($\varepsilon = 0.5$), close to the first order of Lemma 7. A weighted straight-line fit through $\Delta t \leq 0.025$ extrapolates to $V(0, 1) = 1.9098 \pm 0.0021$ at $\varepsilon = 0.05$ and 1.9056 ± 0.0021 at $\varepsilon = 0.5$, against the closed form 1.9077. The two truncation levels share the limit, as the LQ structure predicts (the ε -term of Proposition 3 vanishes); at a fixed Δt they differ because the large-jump clock shortens the sub-steps — 137 against 101 sub-steps per path at $\Delta t = 0.05$ — not because of the Gaussian approximation.

The total error after N paths is about $(\sigma_J^2/N + \text{bias}^2)^{1/2}$. At $\alpha = 1.7$ and $\varepsilon = \Delta t = 0.05$ the bias equals the statistical error at $N = (\sigma_J/\text{bias})^2 \approx 1769$ paths; beyond that, more paths do not help without a smaller Δt .




Figure 3: Monte Carlo statistical error and bias.

Figure 3. Left: root-mean-square statistical error of $\widehat{V}_N(0, x_{\text{ref}})$, relative to $V(0, 1)$, against the path count N for four values of α ($\varepsilon = \Delta t = 0.05$; 95 % intervals), with the reference slope $N^{-1/2}$; the legend gives the fitted slopes with one standard error. Right: bias of the Monte Carlo mean against the closed form as a function of the outer step Δt at $\alpha = 1.7$ for $\varepsilon = 0.05$ and $\varepsilon = 0.5$ (10^6 paths per point, 95 % intervals), with the straight-line fits through $\Delta t \leq 0.025$.



9.8 Sensitivity to α

The optimal feedback gain $u^*(0, x)/x = -(b/\alpha_c) A(0) = -0.8828$ does not depend on α in the LQ subcase, because the non-local term enters only the constant $C(t)$, not the slope $A(t)$. Table 2 shows where α does enter: the truncated second moment, the offset $C(0)$ and the value, and the two Monte Carlo constants — the path-cost spread σ_J and the work per path. This separation is consistent with the verification theorem: the optimal feedback depends on $\partial_x V$, which sees only the quadratic coefficient, while the level of the value function sees the Lévy structure.

Table 2: Where α enters, at $A(0) = 0.8828$ for every α . σ_R^2 , $C(0)$ and $V(0, 1)$ from the closed form; σ_J (the standard deviation of one path's cost), the large-jump rate λ_ε and the mean number of sub-steps per path from the Monte Carlo runs of Figure 3 ($\varepsilon = \Delta t = 0.05$).

α	σ_R^2	$C(0)$	$V(0, 1)$	σ_J	λ_ε	sub-steps per path
1.3	1.1846	2.3399	2.7813	2.959	10.13	128
1.5	0.9462	1.8690	2.3103	2.345	12.60	135
1.7	0.7423	1.4663	1.9077	1.744	13.16	137
1.9	0.5728	1.1314	1.5728	1.117	7.61	121

9.9 Brownian recovery

Proposition 4 can be tested on $\varphi = \cos$, for which $\mathcal{J}_\alpha \cos = -\psi_R(\sigma) \cos$ with ψ_R the truncated exponent of Section 6. The proposition says that $\psi_R(\sigma) \rightarrow \frac{1}{2}\sigma_2^2\sigma^2 = \frac{1}{4}$ at $\sigma = 1$, at rate $O(2 - \alpha)$. Theorem 6 says $V_\alpha \rightarrow V_2$, where V_2 solves the Brownian HJB with diffusion term $\frac{1}{2}\sigma_2^2\sigma^2\partial_x^2$; in the LQ subcase $V_2(0, 1) = 1.4290$. Table 3 computes both by quadrature: the ratios in the fourth and sixth columns settle at constants, so both gaps close at the rate $2 - \alpha$.



Table 3: Proposition 4 and Theorem 6 in the LQ subcase ($R = 5$, $\sigma = 1$). $\psi_R(1)$ by quadrature of $\int_{|z| \leq R} (1 - \cos z) \nu_\alpha(dz)$; $V_\alpha(0, 1)$ from the closed form; $V_2(0, 1) = 1.4290$ from the Brownian Riccati system with $\sigma_2^2 = \frac{1}{2}$.

α	σ_R^2	$\psi_R(1)$	$(\psi_R(1) - \frac{1}{4})/(2 - \alpha)$	$V_\alpha(0, 1)$	$(V_\alpha - V_2)/(2 - \alpha)$
1.3	1.1846	0.38642	0.1949	2.7813	1.9318
1.5	0.9462	0.34419	0.1884	2.3103	1.7626
1.7	0.7423	0.30405	0.1802	1.9077	1.5954
1.9	0.5728	0.26712	0.1712	1.5728	1.4373
1.99	0.5069	0.25167	0.1670	1.4427	1.3700
1.999	0.5007	0.25017	0.1666	1.4304	1.3635

10. Discussion

The verification theorem of Section 5 closes the loop between the viscosity-solution characterisation of the value function under α -stable Lévy noise and the candidate feedback law obtained by pointwise minimisation of the Hamiltonian. Three observations are worth recording.

10.1 Why α matters

Heavy-tailed driving noise increases the running cost via the constant offset $C(t)$ of the value function (Section 9), but does not by itself distort the optimal feedback slope: in the LQ subcase the slope is a function of $(\rho, \theta, b, \alpha_c, \beta, \gamma_T)$ only, and the α -stable Lévy measure enters only through its truncated second moment. The qualitative shape of the optimal control is therefore robust to the choice of α within the truncated regime. Outside the LQ subcase, where the non-local term cannot be collapsed to a constant, the α -dependence of u^* is genuine and is encoded in the resolvent representation of Theorem 5.

10.2 Limitations

Pure α -stable only. The framework is restricted to symmetric truncated α -stable noise. Asymmetric Lévy measures (CGMY, tempered stable, Variance-Gamma, NIG) preserve the non-local structure but introduce a drift–noise dependence that breaks the symmetric simplification of Remark 2. The verification theorem extends to these cases — only the generator definition changes — but the closed-form simplification of the LQ subcase does not.

One-dimensional state. The state is scalar. The viscosity-solution framework of Section 4 extends to higher dimensions; only the doubling-the-variables argument in Appendix A becomes notationally heavier. The Asmussen–Rosiński scheme generalises to multivariate α -stable measures via spherical decomposition.



Truncated Lévy measure. Section 2 truncates the Lévy measure at $|z| \leq R$ to ensure finite second moment. Removing the truncation forces a move to fractional moments $p < \alpha$ for the cost functional, which re-opens the polynomial-growth hypothesis in Theorem 4. The viscosity- solution argument of Section 4 carries through unchanged with the larger test-function class.

Bounded control set. The admissible set U is bounded; this ensures the candidate feedback law from pointwise minimisation always lies in U on the policy-relevant region of state space. Unbounded controls require an additional integrability hypothesis on the optimal feedback along the closed-loop trajectory; the standard fix is a quadratic-growth penalty.

10.3 Extensions

The framework extends naturally in three directions. First, **multi- dimensional state**: the viscosity-solution programme and the verification theorem generalise without difficulty, only the comparison-principle proof in Appendix A becomes notationally heavier. Second, **anisotropic α -stable noise**: the Lévy measure is replaced by a spherical decomposition $\nu(dz) = r^{-1-\alpha} dr \otimes \Sigma(d\omega)$ on $\mathbb{R}^d \setminus \{0\}$; the resolvent representation of Section 6 generalises with the symbol $\sigma^\alpha |\xi|^\alpha$ replaced by the anisotropic Lévy exponent. Third, **path-dependent costs**: replacing the running cost $\frac{1}{2}\beta x^2$ by a path-dependent functional of the realised state introduces a memory kernel; the natural framework is BSDE-based verification, which we do not pursue here.

The most immediate next step is the asymmetric jump-diffusion case (CGMY or tempered stable), where the closed-form LQ subcase admits a similar truncation and the resolvent representation of Section 6 gives a direct algorithmic path to the optimal feedback via Fourier inversion of the characteristic exponent.

11. Conclusion

We have established a verification theorem for finite-horizon Hamilton–Jacobi–Bellman equations driven by symmetric truncated α -stable Lévy noise, in the viscosity-solution framework. A smooth viscosity solution of polynomial growth coincides with the value function, and the feedback law obtained by pointwise minimisation of the Hamiltonian is optimal in the admissible class. The non-local generator is handled by the Barles–Imbert maximum-principle technology; the comparison principle is proven in Appendix A; the verification proof uses the Itô–Lévy formula and a Burkholder–Davis–Gundy bound on the compensated Poisson random measure.

A jump-adapted Monte Carlo scheme for the controlled state, based on the Asmussen–Rosiński small-jump Gaussian approximation, was checked against the closed form of a linear–quadratic truncated- α -stable subcase. Its statistical error decays as $N^{-1/2}$ at every $\alpha \in \{1.3, 1.5, 1.7, 1.9\}$ — the index changes the constant, not the exponent — and its bias is first order in the time step; extrapolated to $\Delta t \rightarrow 0$ it returns the closed-form value within one standard error. The same simulations confirm the verification theorem



(every other linear feedback costs more) and the corrected resolvent representation of the feedback gain.

As $\alpha \rightarrow 2^-$, the non-local generator converges to the local second-order operator and the verification theorem reduces to the classical Fleming–Soner statement. The framework extends naturally to multi-dimensional state, anisotropic Lévy measures, and CGMY / tempered- stable noise via the resolvent representation of Theorem 5, which we leave to subsequent work.

A. Comparison principle: full proof

We prove Theorem 2 by doubling the variables (Crandall–Ishii [6]) with the non-local maximum principle of Jakobsen–Karlsen [12] and the Barles–Imbert small-/large-jump split [4].

Assumption 2 (Hypotheses recalled). $\underline{V}$ is an upper semicontinuous viscosity subsolution and $\bar{V}$ is a lower semicontinuous viscosity supersolution of the reduced HJB PIDE of Theorem 1 on $[0, T] \times \mathbb{R}$. Both functions have polynomial growth: there exist $K > 0, k \geq 0$ such that $|\underline{V}(t, x)| + |\bar{V}(t, x)| \leq K(1 + |x|^k)$. The terminal data satisfy $\underline{V}(T, \cdot) \leq \bar{V}(T, \cdot)$. Assumption 1 holds.

A.1 Step 1. Auxiliary function and contradiction setup

Lemma 1 (Contradiction setup). Suppose, for contradiction, that there exists $(t^*, x^*) \in [0, T] \times \mathbb{R}$ with $\underline{V}(t^*, x^*) > \bar{V}(t^*, x^*)$. Then there exist $\delta_0 > 0$ and a compact $K^* \subset \mathbb{R}$ such that

$$M_0 := \sup_{(t,x) \in [0,T] \times K^*} (\underline{V}(t, x) - \bar{V}(t, x)) > \delta_0 > 0.$$

The supremum is attained at some $(\hat{t}, \hat{x}) \in [0, T] \times K^*$ with $\hat{t} < T$, by upper semicontinuity of $\underline{V} - \bar{V}$ and compactness of K^* .

Proof. Polynomial growth localises the supremum to a compact set; semicontinuity gives the attainment. The strict inequality $\hat{t} < T$ follows from the terminal-data hypothesis. $\square$

For $\eta, \delta > 0$ define on $[0, T]^2 \times K^* \times K^*$

$$\Phi_{\eta, \delta}(t, s, x, y) := \underline{V}(t, x) - \bar{V}(s, y) - \frac{|x - y|^2 + (t - s)^2}{2\eta} - \delta(t + s), \quad (\text{A.1})$$

where the η -penalty pulls the maximum onto the diagonal and the δ -penalty keeps it away from $t = T$.

Lemma 2 (Maximum approaches the diagonal). For $\delta > 0$ fixed and $\eta \rightarrow 0^+$, the maximum $M_{\eta, \delta} := \sup \Phi_{\eta, \delta}$ is attained at a point $(t_\eta, s_\eta, x_\eta, y_\eta)$ with $|x_\eta - y_\eta|^2 / \eta + (t_\eta - s_\eta)^2 / \eta \rightarrow 0$ as $\eta \rightarrow 0^+$, and the limit point lies on the diagonal of $[0, T] \times K^*$.

Proof. Standard penalisation argument; see [6], Section 3, with the additional δ term handled by reducing δ at the end. $\square$



A.2 Step 2. Non-local maximum principle

Lemma 3 (Jakobsen–Karlsen non-local maximum principle). At the maximum point $(t_\eta, s_\eta, x_\eta, y_\eta)$, there exist real numbers a, b with $a - b = \delta$ and scalars $X, Y \in \mathbb{R}$ with $X \leq Y$ such that

$$(a, p_\eta, X) \in \overline{\mathcal{P}}_{\text{NL}}^{2,+} \underline{V}(t_\eta, x_\eta), \quad (b, p_\eta, Y) \in \overline{\mathcal{P}}_{\text{NL}}^{2,-} \overline{V}(s_\eta, y_\eta),$$

where $p_\eta = (x_\eta - y_\eta)/\eta$ and $\overline{\mathcal{P}}_{\text{NL}}^{2,\pm}$ denote the parabolic super-/subjet adapted to the non-local operator $\mathcal{J}$ (Definition 3).

Proof. This is Theorem 3.1 of [12], adapted to our reduced HJB PIDE. The non-local sub-/superjets carry the test-function contribution to $\mathcal{J}$ on the small-jump side and the candidate-function contribution on the large-jump side, following the Barles–Imbert split [4]. $\square$

A.3 Step 3. Estimate at the maximum point

Substituting the sub-/superjet relations into the viscosity definitions and subtracting,

$$\delta = a - b \leq \rho (\overline{V}(s_\eta, y_\eta) - \underline{V}(t_\eta, x_\eta)) + \frac{b^2}{2\alpha_c} p_\eta^2 - \theta (x_\eta - y_\eta) p_\eta + \frac{1}{2} \beta (x_\eta^2 - y_\eta^2) + \mathcal{J}^\eta(t_\eta, s_\eta), \quad (\text{A.2})$$

where $\mathcal{J}^\eta$ is the non-local difference, split à la Barles–Imbert into small-jump (test-function-controlled) and large-jump (candidate-function-controlled) parts; both are $O(\eta) + O(\delta)$.

A.4 Step 4. Passing to the limit and contradiction

As $\eta \rightarrow 0^+$, $(t_\eta, x_\eta), (s_\eta, y_\eta) \rightarrow (\hat{t}, \hat{x})$ on the diagonal; the right-hand side tends to ρM_0 , so

$$\delta \leq \rho M_0 - 2\delta \hat{t}, \quad (\text{A.3})$$

which contradicts $\delta > 0$ for M_0 fixed and δ small.

Theorem 7 (Comparison, restated). Under Assumption 2, $\underline{V} \leq \overline{V}$ on $[0, T] \times \mathbb{R}$.

Proof. Steps 1–4 — Lemma 1, Lemma 2 and Lemma 3, then the estimate and the limit — produce a contradiction from the assumption that there exists a point at which $\underline{V} > \overline{V}$. Hence no such point exists. $\square$

Remark 11 (Polynomial-growth restriction). The polynomial-growth bound is what allows the localisation of Step 1 to a compact set; without it, the maximum can escape to infinity and the doubling argument breaks down. The hypothesis is the natural one for HJB equations driven by Lévy noise — the value function inherits polynomial growth from the cost functional under any bounded admissible control.



B. Convergence analysis of the Monte Carlo scheme

We prove Proposition 3 in full. The error of the empirical value $\widehat{V}_{N,\varepsilon,\Delta t}$ decomposes into three additive contributions; each is bounded separately by a standard argument from the literature, then combined.

B.1 Error decomposition

Lemma 4 (Additive decomposition). Let $V(0, x)$ denote the true value at $(0, x)$, $V_{\varepsilon,\Delta t}(0, x) := \mathbb{E}[\widehat{J}^{\varepsilon,\Delta t}]$ the expected cost under the discretised closed-loop dynamics with small-jump truncation ε and outer step Δt , and $\widehat{V}_N(0, x) := \frac{1}{N} \sum_{m=1}^N \widehat{J}^{(m)}$ the empirical value over N i.i.d. trajectories. Then

$$\mathbb{E}|\widehat{V}_N - V(0, x)| \leq \underbrace{\mathbb{E}|\widehat{V}_N - V_{\varepsilon,\Delta t}|}_{\text{Monte Carlo}} + \underbrace{|V_{\varepsilon,\Delta t} - V_{\Delta t}|}_{\text{small-jump approx.}} + \underbrace{|V_{\Delta t} - V(0, x)|}_{\text{Euler discretisation}}.$$

Proof. Triangle inequality, after inserting the two auxiliary value functions $V_{\Delta t}$ (exact Lévy-driven dynamics with Euler-discretised drift) and $V_{\varepsilon,\Delta t}$ (small-jump Gaussian-approximated Lévy dynamics with Euler drift). $\square$

B.2 Step 1. Monte Carlo

Lemma 5 (Monte Carlo error). Under Assumption 1 and bounded admissible controls, the discretised cost $\widehat{J}^{\varepsilon,\Delta t}$ has finite variance $\sigma_J^2 < \infty$, and

$$\mathbb{E}|\widehat{V}_N - V_{\varepsilon,\Delta t}| \leq \frac{\sigma_J}{\sqrt{N}}.$$

Proof. Independent trajectories yield i.i.d. costs. The cost is quadratic in the state, so its variance is finite as soon as the state has a finite fourth moment, which holds because the truncated jumps have moments of every order and the controls are bounded (the argument of Proposition 1 with fourth powers). By the Cauchy–Schwarz inequality, $\mathbb{E}|\widehat{V}_N - V_{\varepsilon,\Delta t}| \leq (\text{Var } \widehat{V}_N)^{1/2} = \sigma_J/\sqrt{N}$. $\square$

B.3 Step 2. Small-jump Gaussian approximation

Lemma 6 (Asmussen–Rosiński). For each $\varepsilon \in (0, R)$, the difference between the exact Lévy-driven dynamics and the small-jump Gaussian-approximated dynamics satisfies

$$|V_{\varepsilon,\Delta t} - V_{\Delta t}| \leq C_{\text{AR}} \varepsilon^{3-\alpha},$$

for cost functionals whose dependence on the state has three derivatives of polynomial growth, where $C_{\text{AR}} > 0$ depends on (α, σ, T) but not on $(\varepsilon, \Delta t, N)$. For the LQ cost the difference is zero.

Proof. The small-jump Lévy increment over a sub-step is replaced by a Gaussian with the same mean and variance, as in [9]. A second-order Taylor expansion of the cost-



to-go around the increment cancels the terms of order zero, one and two, because the two increments share their first two moments; the remainder is bounded by the third absolute moments, and summing over the sub-steps of $[0, T]$ leaves $C \int_{|z| \leq \varepsilon} |z|^3 \nu_\alpha(dz) = 2C c_\alpha \varepsilon^{3-\alpha} / (3-\alpha)$. For the LQ subcase the cost-to-go is quadratic, the remainder vanishes identically, and the expected cost depends on the noise only through its variance. $\square$

B.4 Step 3. Euler drift discretisation

Lemma 7 (Euler discretisation error). Under Assumption 1 and the regularity $\partial_x V \in L^\infty([0, T] \times \mathbb{R})$ of Proposition 3,

$$|V_{\Delta t} - V(0, x)| \leq C_E \Delta t,$$

where $C_E > 0$ depends on $(\alpha, \sigma, b, \theta, R, T)$ and on $\|\partial_x V\|_\infty$ but not on $(\varepsilon, \Delta t, N)$.

Proof. Standard Euler-scheme analysis for jump-diffusions; see [13], Chapter 6. The Lipschitz drift $bu^*(t, x) - \theta x$ under bounded u^* and linear-in- x growth ensures the per-step error is $O(\Delta t^2)$; accumulation over $T/\Delta t$ steps gives $O(\Delta t)$ in the value functional, and the left-point accumulation of the running cost adds another $O(\Delta t)$. $\square$

B.5 Combining the bounds

Theorem 8 (Convergence, restated). Combining Lemmas 5–7 with the decomposition of Lemma 4,

$$\mathbb{E}|\widehat{V}_N - V(0, x)| \leq \frac{\sigma_J}{\sqrt{N}} + C_E \Delta t + C_{AR} \varepsilon^{3-\alpha}.$$

The statistical term is the only one that falls with N , and it falls as $N^{-1/2}$ for every α . Choosing $\Delta t \asymp N^{-1/2}$ and $\varepsilon^{3-\alpha} \asymp N^{-1/2}$ makes all three terms of the same order, so the total error is $O(N^{-1/2})$; no choice of ε and Δt makes it smaller, since the statistical term does not depend on them.

Proof. Lemma 4 with the three bounds of Lemma 5, Lemma 6 and Lemma 7. The work per path grows like $T/\Delta t + \lambda_\varepsilon T$ with $\lambda_\varepsilon = O(\varepsilon^{-\alpha})$, so the choices above cost $O(N^{1/2} + N^{\alpha/(2(3-\alpha))})$ sub-steps per path. $\square$

Remark 12 (Variance reduction). The constant σ_J grows as α decreases at fixed R , with the truncated second moment σ_R^2 : in Section 9 it is 2.96 at $\alpha = 1.3$ and 1.12 at $\alpha = 1.9$. Variance reduction (a control variate built on the LQ closed form, for instance) lowers the constant; it does not change the exponent $1/2$. We have not measured its gain here.

C. Algorithm pseudocode in full

Three algorithms, given in full pseudocode form. The reference Python implementation is in `hjb_stable.py` (closed forms), `mc_jump_adapted.py` (the simulation), `checks.py` (every number in Section 9) and `contours.py` (the figures), distributed with this preprint.


Algorithm 3 — Closed-form value function — LQ truncated- α -stable subcase

Möbius-interpolation closed form for the Riccati coefficient $A(t)$, adaptive quadrature for the offset $C(t)$.

```

1 function riccati_value(alpha, sigma, theta, b, rho, alpha_c, beta,
  gamma_T, R, T, t, x):
2   # 1. Pre-compute alpha-stable normalisation and truncated second
  moment
3   c_alpha  <- alpha * Gamma((1+alpha)/2) /
  (2*sqrt(pi)*Gamma(1-alpha/2))
4   sigma2_R <- 2*c_alpha * R^(2 - alpha) / (2 - alpha)
5
6   # 2. Riccati roots and decay rate
7   a2 <- b^2 / alpha_c
8   a1 <- rho + 2*theta
9   lambda  <- sqrt(a1^2 + 4*a2*beta)
10  A_plus   <- (-a1 + lambda) / (2*a2)
11  A_minus  <- (-a1 - lambda) / (2*a2)
12
13  # 3. Mobius closed form for A(t), with A(T) = gamma_T
14  s  <- T - t
15  e  <- exp(-lambda*s)
16  num <- A_plus*(gamma_T - A_minus) - A_minus*(gamma_T - A_plus)*e
17  den <- (gamma_T - A_minus) - (gamma_T - A_plus)*e
18  A_t <- num / den
19
20  # 4. C(t) = 0.5*sigma^2*sigma2_R * int_0^s exp(-rho*(s - r)) *
  A_of_s(r) dr,
21  #   with A_of_s(r) the Mobius formula of step 3 at time-to-horizon r
22  C_t <- 0.5*sigma^2*sigma2_R * adaptive_quadrature(r -> exp(-rho*(s -
  r))*A_of_s(r), 0, s)
23
24  # 5. Quadratic value function
25  return 0.5*A_t*x^2 + C_t

```

Algorithm 4 — Jump-adapted single-step transition

Asmussen-Rosiński small-jump Gaussian + large-jump Poisson clock between drift integration steps.

```

1 function jump_adapted_step(params, x, t, eps, dt_max, T, rng, feedback):
2   # Small-jump Gaussian variance and large-jump rate
3   c_alpha  <- stable_density_const(alpha)
4   sigma_eps <- sqrt(2*c_alpha * eps^(2 - alpha) / (2 - alpha))
5   lambda_eps <- 2*c_alpha / alpha * (eps^(-alpha) - R^(-alpha))
6
7   # Next large-jump waiting time
8   dt_jump <- rng.exponential(1 / lambda_eps)      if lambda_eps > 0
  else inf

```



```

9     step    <- min(dt_max, dt_jump, T - t)
10
11     # Optimal control evaluated at current state
12     u_star  <- feedback(t, x)                # =
13     -(b/alpha_c)*A(t)*x
14
15     # Drift + small-jump Gaussian increment over [t, t + step]
16     x_pre   <- x + (b*u_star - theta*x)*step
17             + sigma * sigma_eps * sqrt(step) * rng.normal()
18
19     # Did the large-jump clock ring inside the step, before T?
20     if dt_jump < min(dt_max, T - t):
21         u <- rng.uniform()
22         sg <- +1 if rng.uniform() < 0.5 else -1
23         base <- eps^(-alpha) - u * (eps^(-alpha) - R^(-alpha))
24         z <- sg * base^(-1/alpha)                # truncated-stable
25         jump
26         x_post <- x_pre + sigma * z
27     else:
28         x_post <- x_pre
29
30     return (x_post, t + step)

```

Algorithm 5 — Full trajectory simulator under optimal feedback

Iterated calls to Algorithm 4 until horizon T ; accumulates the discounted running cost and the terminal cost for a single path.

```

1 function simulate_path(params, x0, eps, dt_max, T, rng, feedback):
2     t <- 0
3     x <- x0
4     J <- 0
5     while t < T:
6         u_star <- feedback(t, x)
7         # Discounted instantaneous cost rate (recorded at sub-step start)
8         L_rate <- 0.5*alpha_c*u_star^2 + 0.5*beta*x^2
9         (x_new, t_new) <- jump_adapted_step(params, x, t, eps, dt_max, T,
10 rng, feedback)
11         # Left-point accumulation of the running cost over [t, t_new]
12         J <- J + exp(-rho*t) * L_rate * (t_new - t)
13         x, t <- x_new, t_new
14     J <- J + exp(-rho*T) * 0.5*gamma_T*x^2                # terminal cost
15     return J
16
17 function empirical_value(params, x0, N, eps, dt_max, T, seed, feedback):
18     rng <- rng_init(seed)
19     J_paths[m] <- simulate_path(params, x0, eps, dt_max, T, rng, feedback)
20     for m = 1..N

```



```

20     V_hat <- mean(J_paths)
21     se     <- std(J_paths) / sqrt(N)
22     return (V_hat, se)

```

Remark 13 (Pipeline). Algorithm 3 produces the closed-form value; Algorithm 4 is the inner kernel of the jump-adapted Monte Carlo; Algorithm 5 wraps Algorithm 4 with the cost accumulation along a path. The reference implementation also carries the pathwise derivative Φ of Theorem 5 along each path and evaluates the cost of the feedback $u = -m(b/\alpha_c)A(t)x$ for any m . Section 9 compares the output with Algorithm 3: the statistical error against N and the bias against Δt (Figure 3), and the costs of five feedbacks (Table 1).

Remark 14 (Reproducibility). The simulations take an explicit seed, and `checks.py` records every seed it uses next to the numbers it writes; `contours.py` draws the figures from that record. Re-running with the same seeds reproduces the reported numbers up to platform-specific floating-point order of operations.



References

- [1] Fleming, W. H., & Soner, H. M. (2006). *Controlled Markov Processes and Viscosity Solutions*. Springer.
- [2] Yong, J., & Zhou, X. Y. (1999). *Stochastic Controls: Hamiltonian Systems and HJB Equations*. Springer.
- [3] Pham, H. (1998). Optimal stopping of controlled jump diffusion processes: a viscosity solution approach. *Journal of Mathematical Systems, Estimation and Control*, 8, 1–27.
- [4] 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>
- [5] Mikulevičius, R., & Pragarauskas, H. (2014). On the Cauchy problem for integro-differential operators in Sobolev classes and the martingale problem. *Journal of Differential Equations*, 256(4), 1581–1626.
<https://doi.org/10.1016/j.jde.2013.11.008>
- [6] 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>
- [7] 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>
- [8] Levendorskiĭ, S. Z. (2004). Pricing of the American put under Lévy processes. *International Journal of Theoretical and Applied Finance*, 7(3), 303–335.
<https://doi.org/10.1142/s0219024904002463>
- [9] Asmussen, S., & Rosiński, J. (2001). Approximations of small jumps of Lévy processes with a view towards simulation. *Journal of Applied Probability*, 38(2), 482–493. <https://doi.org/10.1239/jap/996986757>
- [10] Sato, K.-i. (1999). *Lévy Processes and Infinitely Divisible Distributions*. Cambridge University Press.
- [11] Applebaum, D. (2009). *Lévy Processes and Stochastic Calculus*. Cambridge University Press. <https://doi.org/10.1017/cbo9780511809781>
- [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 NoDEA*, 13(2), 137–165.
<https://doi.org/10.1007/s00030-005-0031-6>
- [13] Glasserman, P. (2003). Monte Carlo Methods in Financial Engineering. *Springer Stochastic Modelling and Applied Probability*, 53.



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).





Phase portrait of the Lotka–Volterra equations: $\dot{x} = \alpha x - \beta xy$, $\dot{y} = -\gamma y + \delta xy$, $\alpha = 1.09$, $\beta = 0.71$, $\gamma = 1.20$, $\delta = 1.00$



ZAMRIK QUANT RESEARCH

The research arm of Zamrik Quant Lab



PR-2026-22911048
zamrik.com