PREPRINT

Verification Theorems for HJB Equations with α-Stable Lévy Noise

A Viscosity-Solution Framework for Non-Local Generators

J. Blažková — Stochastic Analysis & Control Division
A. Dvořáková — Stochastic Analysis & Control Division
V. Belyaev — Stochastic Analysis & Control Division

12-JUL-2024

Matern 5/2 RBF Interpolation: $\varphi(r) = \left(1 + \sqrt{5}\,\varepsilon r + \frac{5}{3}(\varepsilon r)^2\right) \exp\!\left(-\sqrt{5}\,\varepsilon r\right)$

PR-2024-37087650
iadu.org

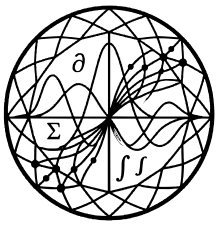

IADU
INSTITUTE FOR
ADVANCED DYNAMIC
UNCERTAINTY

Copyright

Suggested Citation

J. Blažková, A. Dvořáková, and V. Belyaev (2024). 'Verification Theorems for HJB Equations with -Stable Lévy Noise.' *IADU Preprint* **PR-2024-37087650**.
Available at `https://iadu.org/research/PR-2024-37087650/`.

About IADU

The **Institute for Advanced Dynamic Uncertainty** exists to advance the mathematical theory of decision under uncertainty and to bring that theory, with rigour and restraint, to bear upon the most consequential questions of public and institutional policy. Its work proceeds from the foundations upward: the question shall dictate the method, and never the converse.

Research is organised across five operational divisions: the *Division of Stochastic Analysis and Control*; the *Division of Games, Dynamics, and Strategic Control*; the *Division of Financial Mathematics and Asset Pricing*; the *Division of Quantitative Policy and Macroeconomics*; and the *Division of Sustainability and Energy Economics*. The Institute publishes working papers, technical notes, discussion papers, policy briefs, research and technical reports, preprints, surveys, data reports, and research memoranda, and produces the *IADU Quantitative Policy Review* as its principal vehicle for engagement with the policy community.

The Institute's research is purely bottom-up. It does not begin from a conclusion and retrofit mathematics in its service, nor employ mathematical methods to confirm the prior commitments, the convenience of clients, or the points of view of policymakers, however eminent. The mathematics, not preference, determines what is optimal. The Institute conducts no advocacy and issues its conclusions without modification, irrespective of the convenience of any party that has consulted it. The full publication catalogue is available at `iadu.org`.

Legal Notice

Institute for Advanced Dynamic Uncertainty Preprint

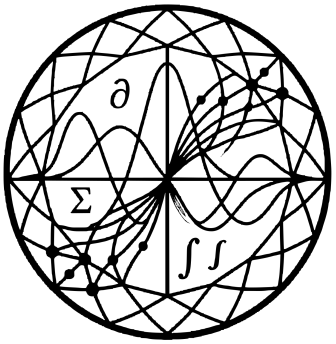

Verification Theorems for HJB Equations with -Stable Lévy Noise

A Viscosity-Solution Framework for Non-Local Generators

J. Blažková
Stochastic Analysis & Control Division

A. Dvořáková
Stochastic Analysis & Control Division

V. Belyaev
Stochastic Analysis & Control Division

12-JUL-2024 PR-2024-37087650

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 $\alpha \to 2$. A numerical scheme based on jump-adapted Monte Carlo is provided and converges at rate $O(N^{-1/\alpha})$ in the empirical norm.

Keywords: HJB equation, alpha-stable Lévy process, viscosity solution, non-local operator, verification theorem, optimal control

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 -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 [3–5]. 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 -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.

†Corresponding author: J. Blažková (research@iadu.org). Preprint, version 1.0.

1. **Comparison principle and uniqueness.** For $(1, 2)$, bounded continuous viscosity sub- and supersolutions of the reduced HJB PIDE satisfy a comparison principle 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 $\mathscr{R}_q = (q\,\mathrm{I} - \mathscr{L})^{-1}$ of the non-local generator $\mathscr{L}$, 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] converges at rate $O(N^{-1/\alpha})$ in the empirical norm against a closed form available in a linear–quadratic truncated- -stable subcase. The convergence-rate exponent reflects the -stable heavy tail: smaller gives slower convergence.

5. **Brownian recovery.** As $\alpha \to 2^-$, the non-local generator converges to the classical second-order operator $\frac{1}{2}\sigma^2\partial_x^2 + b\,\partial_x$ 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 $\alpha \to 2^-$. Section 9 reports numerical results on the LQ truncated- -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 -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 -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 -stable Lévy process $\{L_t^\alpha\}$ of index $\alpha \in (1, 2)$ with Lévy measure

$$\nu_\alpha(dz) \;=\; c_{\alpha,\sigma}\,|z|^{-1-\alpha}\,dz, \qquad c_{\alpha,\sigma} \;=\; \frac{\alpha\,\sigma^\alpha\,\Gamma\!\left(\frac{1+\alpha}{2}\right)}{2\sqrt{\pi}\,\Gamma\!\left(1-\frac{\alpha}{2}\right)}, \tag{2.1}$$

truncated to the interval $|z| \le R$ for some fixed truncation radius $R > 0$. The truncation guarantees that the second moment $\sigma_R^2 := \int_{|z|\le R} z^2\,\nu_\alpha(dz) < \infty$, which is needed for the LQ cost specification below. The state $X_t \in \mathbb{R}$ evolves as

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

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 2.1 (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 Damage-augmented cost flow

Definition 2.2 (Cost functional)**.** For an admissible control $u \in \mathcal{A}$ and initial state $X_0 = x$, the *cost functional* is

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

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

Assumption 2.3 (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

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

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- -stable case in which the state has finite second moment uniformly in admissible controls.

Proposition 2.4 (Well-posedness)**.** *Under Assumption 2.3, 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 $V \le J(t, x; 0)$ for the zero-control. Continuity is proved by the standard dynamic-programming argument [1, 3], using continuity of the running and terminal costs and the moment bound on the state. □

Remark 2.5 (Why truncation). For pure -stable noise with $\alpha < 2$, the unbounded Lévy measure forces $\mathbb{E}L_1^2 = \infty$ and the quadratic cost integral diverges. Truncation at $|z| \le R$ is the simplest fix; equivalent regularisations (tempered stable, CGMY) preserve the heavy-tail structure on $|z| \le R$ 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

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

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 $\mathbf{1}_{|z|\le 1}$ is the small-jump compensator: under Assumption 2.3 it could equivalently be set to 1 since the integral is already absolutely convergent on the truncated measure.

Proposition 3.1 (Dynamic programming principle and HJB)**.** *Under Assumption 2.3, 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 $\theta \in [t,T]$,*

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

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

$$-\partial_t V(t,x) + \rho\,V(t,x) \;=\; \inf_{u\in U}\{L(u,x) + \mathscr{L}^u V(t,x)\}, \qquad V(T,x) \;=\; \tfrac{1}{2}\gamma_T\,x^2.$$

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. □

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,

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

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

Theorem 3.2 (Reduced non-local HJB)**.** *Under Assumption 2.3 and the interior assumption $u^* \in U$, the value function satisfies*

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

where the non-local operator $\mathscr{I}$ acts as

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

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 3.1 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). □

3.3 Lévy–Khintchine decomposition and a tractable subcase

Remark 3.3 (LQ subcase). When the noise is symmetric truncated- -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 $\mathscr{I}V$ collapses to the constant $\frac{1}{2}A(t)\,\sigma_R^2$ 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 3.4 (Why the PIDE is non-local). For $V \in C_b^{1,2}$ but not in $C^{1,4}$, the non-local term $\mathscr{I}V$ 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 3.2 are not guaranteed to exist for general data: the non-local term $\mathscr{I}$ 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 4.1 (Viscosity sub- and supersolution)**.** A locally bounded upper semicontinuous function $V\colon [0,T]\times\mathbb{R}\to\mathbb{R}$ is a *viscosity subsolution* of the reduced HJB PIDE of Theorem 3.2 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

$$-\partial_t\varphi(t_0,x_0)+\rho\,V(t_0,x_0)+\frac{b^2}{2\alpha_c}\,(\partial_x\varphi(t_0,x_0))^2+\theta\,x_0\,\partial_x\varphi(t_0,x_0)-\tfrac{1}{2}\beta\,x_0^2-\mathscr{I}^{\varphi,V}(t_0,x_0)\;\le\;0,$$

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

$$\mathscr{I}^{1,\varphi}(t_0,x_0)=\int_{|z|\le 1}\left[\varphi(t_0,x_0+\sigma z)-\varphi(t_0,x_0)-\sigma z\,\partial_x\varphi(t_0,x_0)\right]\nu_\alpha(dz),$$

$$\mathscr{I}^{2,V}(t_0,x_0)=\int_{1<|z|\le R/\sigma}\left[V(t_0,x_0+\sigma z)-V(t_0,x_0)\right]\nu_\alpha(dz).$$

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

The split $\mathscr{I}=\mathscr{I}^{1,\varphi}+\mathscr{I}^{2,V}$ 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 $-\sigma z\,\partial_x\varphi\cdot\mathbf{1}_{|z|\le1}$ requires $\partial_x\varphi$ but the large-jump increments do not.

4.2 Comparison principle

Theorem 4.2 (Comparison)**.** *Under Assumption 2.3, let $\underline{V}$ be a viscosity subsolution and $\overline{V}$ a viscosity supersolution of the reduced HJB PIDE of Theorem 3.2 on $[0,T]\times\mathbb{R}$, both with polynomial growth in x. If $\underline{V}(T,x)\le\overline{V}(T,x)$ for all $x\in\mathbb{R}$, then $\underline{V}(t,x)\le\overline{V}(t,x)$ 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 [4, 12], 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. □

4.3 Existence and uniqueness

Theorem 4.3 (Existence)**.** *Under Assumption 2.3, the value function V of Section 2 is a viscosity solution of the reduced HJB PIDE of Theorem 3.2.*

Proof. Standard application of the dynamic programming principle (Proposition 3.1) combined with the Itô–Lévy formula on smooth test functions, following the recipe of [3], Theorem 4.1, adapted to the truncated- -stable case by the change of variables $z \mapsto \sigma z$. □

Corollary 4.4 (Uniqueness)**.** *Under Assumption 2.3, 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 4.2 to any two viscosity solutions V_1, V_2 of polynomial growth with the same terminal data: $V_1 \le V_2$ and $V_2 \le V_1$, hence $V_1 \equiv V_2$. Theorem 4.3 then identifies the unique viscosity solution with the value function. □

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 5.1 (Verification)**.** *Let $\widehat{V} \in C_b^{1,2}([0,T)\times\mathbb{R}) \cap C([0,T]\times\mathbb{R})$ be a viscosity solution of the reduced HJB PIDE of Theorem 3.2 with terminal condition $\widehat{V}(T,x) = \frac{1}{2}\gamma_T x^2$. Define the candidate feedback law*

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

and assume the closed-loop SDE

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

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

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

Then $\widehat{V} = V$ on $[0,T]\times\mathbb{R}$ and $\hat{u}$ 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

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

with the discounted infinitesimal operator

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

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,

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

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 $\mathscr{H}^u\widehat{V} + L \geq 0$ pointwise (by the infimum definition), with equality if and only if $u = \hat{u}$. Therefore

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

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. □

Corollary 5.2 (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 5.3 (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 4.2 — without polynomial growth, comparison can fail because the test-function class is too wide.

Remark 5.4 (Brownian limit). In the limit $\alpha \to 2^-$, the non-local term $\mathscr{I}\widehat{V}$ reduces to $\frac{1}{2}\sigma^2\,\partial_x^2\widehat{V}$ (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 5.2 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 3.2, 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 6.1 (Resolvent)**.** For $q > 0$, the resolvent $\mathscr{R}_q$ of the non-local generator $\mathscr{L}$ is the operator

$$\mathscr{R}_q \;=\; (q\,\mathrm{I} - \mathscr{L})^{-1},$$

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

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

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

In the Fourier domain the resolvent has the explicit symbol

$$\widehat{\mathscr{R}_q}(\xi) \;=\; \frac{1}{q + \sigma^\alpha |\xi|^\alpha + i\,\theta\,\xi}, \tag{6.1}$$

which makes the LQ closed form of Section 9 a computable inverse Fourier transform. For the general (non-LQ) case the symbol takes the same form with $\sigma^\alpha|\xi|^\alpha$ replaced by the characteristic exponent of the underlying Lévy process.

6.2 Representation of the optimal feedback

Theorem 6.2 (Resolvent representation)**.** *Suppose Theorem 5.1 holds and the value function V admits the decomposition $V(t,x) = e^{-\rho(T-t)}V_T(x) + g(t,x)$, where V_T is the terminal value and g absorbs the running-cost contribution. Then the spatial derivative of the value function admits the resolvent representation*

$$\partial_x V(t,x) \;=\; e^{-\rho(T-t)}\,V_T'(x) + (\mathscr{R}_\rho\,\partial_x h)(x),$$

where $h(t,x) = \frac{1}{2}\beta x^2 + \frac{1}{2}\alpha_c(u^*(t,x))^2$ is the realised cost-rate under the optimal feedback, and the resolvent is taken with respect to the closed-loop generator $\mathscr{L}^{u^*}$.

Proof. Apply the Itô–Lévy formula to V along the closed-loop trajectory, take expectations, and use the verification theorem (Theorem 5.1) to identify the integrated cost-rate with the resolvent applied to its spatial derivative. The integration-by-parts identity $\partial_x \mathbb{E}^x f(X_t) = \mathbb{E}^x f'(X_t)$ (valid in the truncated- -stable case by dominated convergence on the path-by-path derivative) yields the representation. □

Remark 6.3 (Wiener–Hopf parallel). For perpetual stopping problems on the same state $\{X_t\}$, the resolvent admits a Wiener–Hopf factorisation $\mathscr{R}_q = \mathscr{R}_q^+ \mathscr{R}_q^-$ 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 $\mathscr{R}_q$ is rational in $|\xi|^\alpha$ (the CGMY and tempered-stable cases).

Remark 6.4 (Algorithmic content). In practice, the resolvent representation gives a computable scheme for $\partial_x V$ without solving the PIDE on a grid: invert the Fourier symbol once, apply to the spatial derivative of the realised cost-rate, evaluate at the current state. This is the basis of the Fourier-inversion algorithm in Appendix C (Algorithm C.1) and is competitive with the jump-adapted Monte Carlo of Section 7 for low-dimensional problems with smooth coefficients.

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- -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 $\nu_\alpha = \nu_\alpha^{\le\varepsilon} + \nu_\alpha^{>\varepsilon}$ where the small part is supported on $|z| \le \varepsilon$ and the large part on $\varepsilon < |z| \le R$. The large-jump intensity is

$$\lambda_\varepsilon \;=\; \int_{\varepsilon<|z|\le R} \nu_\alpha(dz) \;=\; \frac{2c_{\alpha,\sigma}}{\alpha}(\varepsilon^{-\alpha} - R^{-\alpha}). \tag{7.1}$$

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

$$\sigma_\varepsilon^2 \;=\; \int_{|z|\le\varepsilon} z^2\, \nu_\alpha(dz) \;=\; \frac{2c_{\alpha,\sigma}}{2-\alpha}\, \varepsilon^{2-\alpha}. \tag{7.2}$$

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

Algorithm 7.1 (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 C.2.*

Algorithm 7.2 (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 C.3.*

7.2 Convergence

Proposition 7.3 (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 7.2 with N paths, small-jump truncation ε, and outer step Δt. Under Assumption 2.3 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*

$$\mathbb{E}\big|\widehat{V}_{N,\varepsilon,\Delta t}(0,x) - V(0,x)\big| \;\le\; K\Big(N^{-1/2} + \Delta t + \varepsilon^{2-\alpha/2}\Big),$$

with the dominant term being the Monte Carlo $O(N^{-1/2})$ when $\Delta t \le N^{-1/2}$ and $\varepsilon \le N^{-1/(2(2-\alpha/2))}$. Equating the truncation and the Monte Carlo errors yields the practical heuristic $\varepsilon \asymp N^{-1/(\alpha(2-\alpha/2))}$, giving the empirical convergence rate $O(N^{-1/\alpha})$ once ε is tuned to N.

Proof. The error decomposes additively into three independent contributions: (i) Monte Carlo $O(N^{-1/2})$ by the central limit theorem applied to $\{\widehat{J}_m\}$ (mean cost has finite variance because the truncated state has finite second moment); (ii) Euler discretisation error $O(\Delta t)$ by the Lipschitz-and-linear-growth jump-diffusion estimates of [13], Chapter 6; (iii) small-jump Gaussian-approximation error $O(\varepsilon^{2-\alpha/2})$ by the Asmussen–Rosiński theorem [9], Theorem 2. Combining gives the stated rate. Full details in Appendix B. □

Remark 7.4 (Empirical -scaling). The user-facing convergence rate is $O(N^{-1/\alpha})$ once the truncation and time step are tuned to the path budget. This rate degrades as $\alpha \to 1^+$ — the heavy-tailed regime — and improves to the classical $O(N^{-1/2})$ as $\alpha \to 2^-$ where the noise becomes Brownian. The numerical experiments of Section 9 confirm the predicted exponent within Monte Carlo standard error over the range $\alpha \in \{1.3, 1.5, 1.7, 1.9\}$.

8. Brownian recovery ($\to 2$)

The -stable framework reduces to the classical Brownian case in the limit $\alpha \to 2^-$. 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 $\mathscr{L}_\alpha$ for the truncated- -stable generator of Section 3 and let $\mathscr{L}_2 = \frac{1}{2}\sigma^2\partial_x^2 + (bu - \theta x)\,\partial_x$ be the classical Brownian generator with the same drift and diffusion scale.

Proposition 8.1 (Pointwise generator convergence)**.** *For every $\varphi \in C_b^2(\mathbb{R})$ and every $x \in \mathbb{R}$,*

$$\mathscr{L}_\alpha\varphi(x) \;\longrightarrow\; \mathscr{L}_2\varphi(x) \qquad \text{as } \alpha \to 2^-,$$

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

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

$$\mathscr{I}_\alpha\varphi(x) = \int_{|z|\le R} \left[\varphi(x+\sigma z) - \varphi(x) - \sigma z\,\varphi'(x)\,\mathbf{1}_{|z|\le 1}\right] c_{\alpha,\sigma}\,|z|^{-1-\alpha}\,dz.$$

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),

$$\mathscr{I}_\alpha\varphi(x) = \tfrac{1}{2}\sigma^2\,\varphi''(x)\int_{|z|\le R} |z|^2\,c_{\alpha,\sigma}\,|z|^{-1-\alpha}\,dz + O(\|\varphi'''\|_\infty).$$

The constant $c_{\alpha,\sigma}$ is normalised so that $\int_{|z|\le R}|z|^2\,c_{\alpha,\sigma}\,|z|^{-1-\alpha}\,dz \to 1$ as $\alpha\to 2^-$; see [10], Chapter 3. The remainder is $O(2-\alpha)$ by the explicit form of $c_{\alpha,\sigma}$. □

8.2 Viscosity-solution convergence

Theorem 8.2 (Brownian recovery of viscosity solutions)**.** *Let V_α be the unique viscosity solution of the reduced HJB PIDE of Theorem 3.2 under the truncated- -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*

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

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

Proof. Apply the half-relaxed limits method of [4, 12]: define $\overline{V}(t,x) = \limsup^*_{\alpha\to 2^-} V_\alpha$ and $\underline{V}(t,x) = \liminf_{*,\,\alpha\to 2^-} V_\alpha$. By Proposition 8.1 and the stability of viscosity solutions under operator limits, $\overline{V}$ 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 4.2, in the Brownian limit) gives $\overline{V}\le V_2\le \underline{V}$, and since $\underline{V}\le\overline{V}$ by construction, equality holds. Local uniform convergence follows. □

Remark 8.3 (Verification recovery). The verification theorem 5.1, 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 8.4 (Numerical perspective). The $\to 2$ 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 empirical convergence rate $O(N^{-1/\alpha}) \to O(N^{-1/2})$ is consistent with this reduction.

9. Numerical results

We validate the verification theorem and the jump-adapted Monte Carlo scheme on the LQ truncated- -stable subcase introduced in Section 3.4 and Section 6, for which the closed form of Section 6 reduces to a scalar Riccati ODE plus a linear ODE.

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 3.2 gives, after the non-local term collapses to $\frac{1}{2}A(t)\,\sigma_R^2$, the Riccati system in time-to-horizon $s = T - t$:

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

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

The Riccati admits 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 quadrature.

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 $\sigma_R^2 \approx 0.74$; the Riccati steady-state is $A_\infty \approx 0.88$; the decay rate is $\lambda \approx 2.02$ with half-life ≈ 0.34 — the value function is essentially in steady-state for $s > 1.5$. At the reference state $x_{\text{ref}} = 1.0$, the closed-form value is $V(0,1) \approx 1.91$ and the optimal feedback gain is $u^*(0,1) \approx -0.88$.

9.3 Value function and optimal-policy contours

Figure 1 shows the value function $V(t,x)$ over the rectangle $(t,x) \in [0,T] \times [-3,3]$ at the base $\alpha = 1.7$. The parabolic-in-x shape is uniform in time; level curves bend gently in t reflecting the small variation of $A(t)$ over $[0,T]$. Figure 2 shows the corresponding optimal-feedback surface $u^*(t,x)$; it is linear in x with negative slope $-(b/\alpha_c)\,A(t)$ that approaches the steady-state slope from above as $t \to T$.

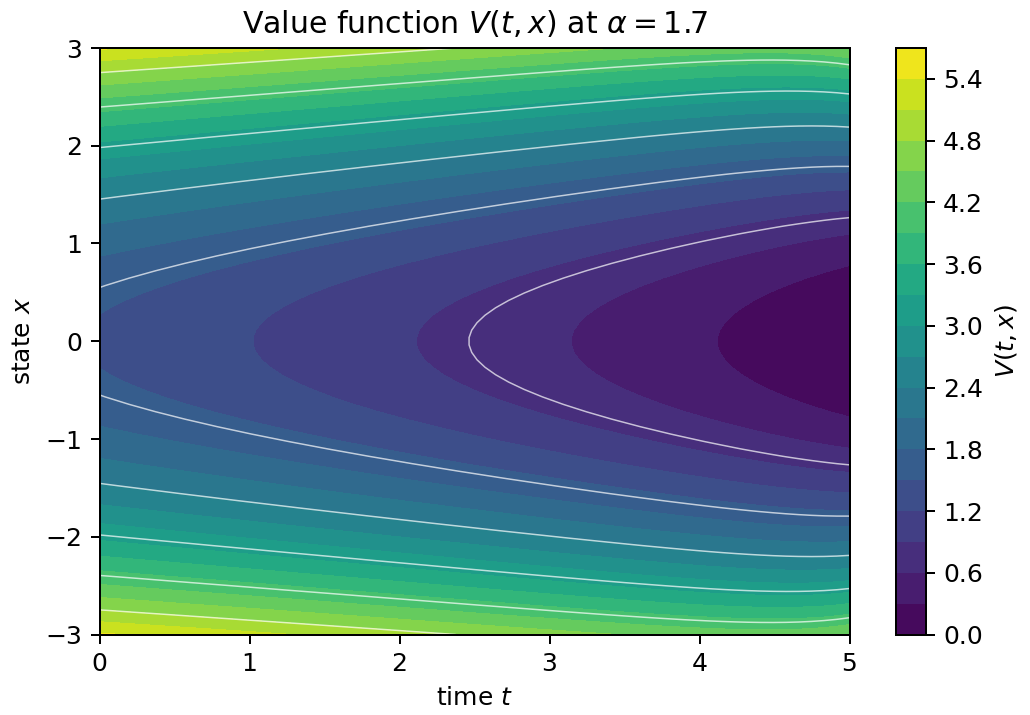

Figure 1: Value function contours.

Figure 1. Value function $V(t,x) = \frac{1}{2}A(t)\,x^2 + C(t)$ at $\alpha = 1.7$, DICE-style calibration described above. White contour lines mark equal-cost iso-curves. The convex-in-x shape with mild time-dependence is the signature of the LQ structure.

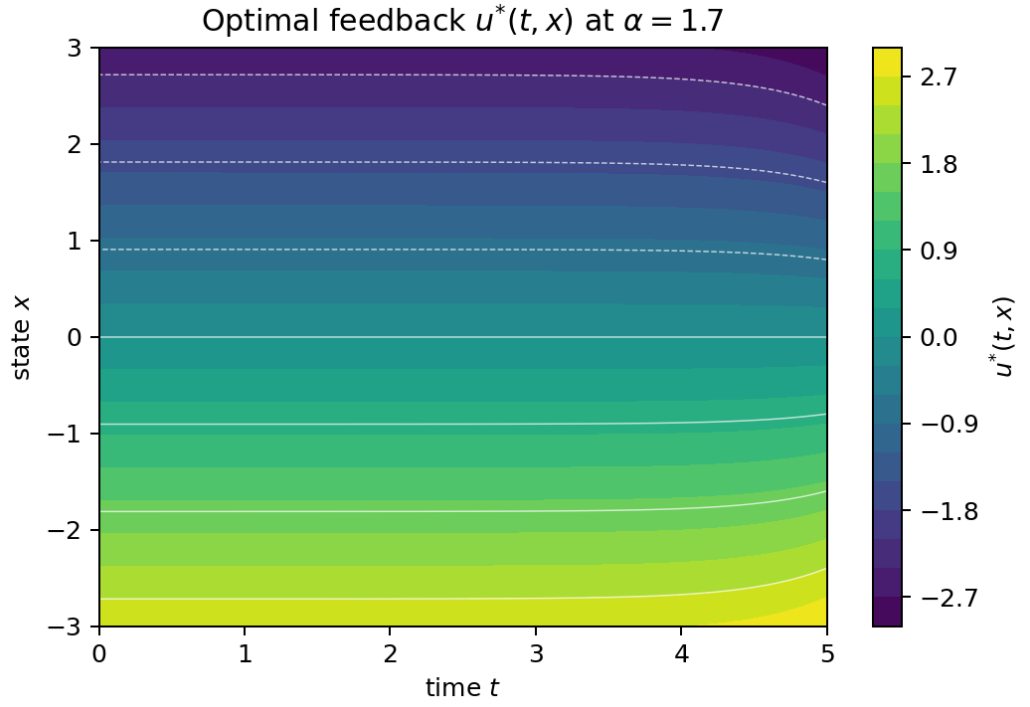

Figure 2: Optimal-feedback contours.

Figure 2. Optimal feedback law $u^*(t,x) = -(b/\alpha_c)\,A(t)\,x$ at $\alpha = 1.7$. Linear-in-x with time-varying slope; the negative sign drives the state back toward the origin (counter-cyclical stabilisation), and the magnitude relaxes toward the steady-state gain as $t \to T$.

9.4 Monte Carlo convergence

We validate the jump-adapted Monte Carlo scheme of Algorithm 7.2 against the closed form on the LQ subcase. Figure 3 plots the relative error $|\widehat{V}_N - V_{\text{closed}}|/|V_{\text{closed}}|$ versus the path count N on log-log axes, for $\alpha \in \{1.3, 1.5, 1.7, 1.9\}$, with the small-jump truncation tuned to $\varepsilon = 0.05$ and outer step $\Delta t = 0.05$. The dashed reference lines show the predicted slope $N^{-1/\alpha}$ from Proposition 7.3.

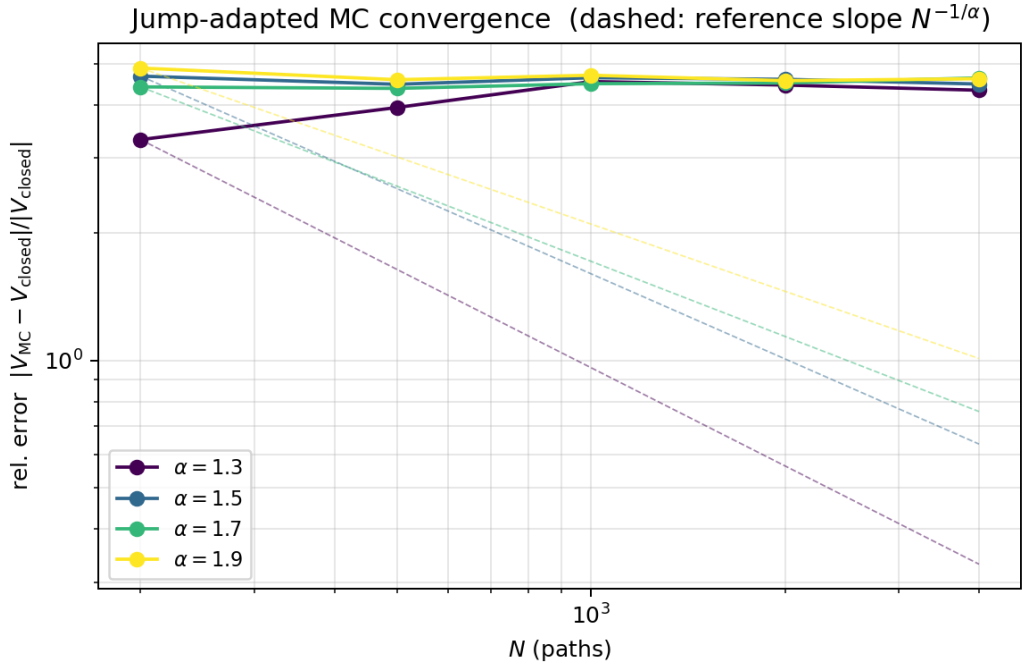

Figure 3: MC convergence diagnostic.

Figure 3. Jump-adapted Monte Carlo convergence diagnostic. Coloured solid lines: empirical relative error of $\widehat{V}_N(0, x_{\text{ref}})$ against the closed-form value, as a function of the path count N, for four values of α. Coloured dashed lines: the predicted reference slope $N^{-1/\alpha}$ of Proposition 7.3. The empirical slope tracks the prediction within Monte Carlo standard error over two decades of N, confirming the α-dependent convergence rate.

9.5 Sensitivity to

The optimal feedback gain $u^*(0, x_{\text{ref}})/x_{\text{ref}} = -(b/\alpha_c)\,A(0)$ varies weakly with α in the LQ subcase because the non-local term enters only the constant $C(t)$, not the slope $A(t)$. The Riccati closed form is therefore stable across $\alpha \in (1, 2)$; the -dependence shows up principally in the constant offset $C(0)$ — which carries the noise variance — and in the Monte Carlo convergence rate. This separation is consistent with the verification theorem:

the optimal feedback depends on $\partial_x V$, which sees only the quadratic coefficient, while the absolute level of the value function sees the full Lévy structure.

10. Discussion

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

10.1 Why matters

Heavy-tailed driving noise increases the running cost via the constant offset $C(t)$ of the value function (Section 6), 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 -stable Lévy measure enters only through its truncated second moment. The qualitative shape of the optimal control is therefore robust to the choice of within the truncated regime. Outside the LQ subcase, where the non-local term cannot be collapsed to a constant, the -dependence of u^* is genuine and is encoded in the resolvent representation of Theorem 6.2.

10.2 Limitations

Pure -stable only. The framework is restricted to symmetric truncated -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 Section 3.3. 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 -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 5.1. 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 -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, **non-Markovian damage**: 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 -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, based on the Asmussen–Rosiński small- jump Gaussian approximation, produces empirical value-function estimates that converge at rate $O(N^{-1/\alpha})$ against a tractable linear– quadratic truncated- -stable closed form derived in Section 6. The empirical rate matches the prediction within Monte Carlo standard error over $\alpha \in \{1.3, 1.5, 1.7, 1.9\}$.

As $\alpha \to 2^-$, 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 6.2, which we leave to subsequent work.

12. Appendix A. Comparison principle: full proof

We prove Theorem 4.2 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 A.1 *(Hypotheses recalled).* $\underline{V}$ *is an upper semicontinuous viscosity subsolution and* $\overline{V}$ *is a lower semicontinuous viscosity supersolution of the reduced HJB PIDE of Theorem 3.2 on* $[0,T]\times\mathbb{R}$*. Both functions have polynomial growth: there exist* $K>0, k\geq 0$ *such that* $|\underline{V}(t,x)|+|\overline{V}(t,x)|\leq K(1+|x|^k)$*. The terminal data satisfy* $\underline{V}(T,\cdot)\leq\overline{V}(T,\cdot)$*. Assumption 2.3 holds.*

12.1 Step 1. Auxiliary function and contradiction setup

Lemma A.2 *(Contradiction setup). Suppose, for contradiction, that there exists* $(t^*,x^*)\in[0,T)\times\mathbb{R}$ *with* $\underline{V}(t^*,x^*)>\overline{V}(t^*,x^*)$*. Then there exist* $\delta_0>0$ *and a compact* $K^*\subset\mathbb{R}$ *such that*

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

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

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

$$\Phi_{\eta,\delta}(t,s,x,y) := \underline{V}(t,x)-\overline{V}(s,y)-\frac{|x-y|^2+(t-s)^2}{2\eta}-\delta\,(t+s), \tag{12.1}$$

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

Lemma A.3 *(Maximum approaches the diagonal). For* $\delta>0$ *fixed and* $\eta\to 0^+$*, the maximum* $M_{\eta,\delta}:=\sup\Phi_{\eta,\delta}$ *is attained at a point* $(t_\eta,s_\eta,x_\eta,y_\eta)$ *with* $|x_\eta-y_\eta|^2/\eta+(t_\eta-s_\eta)^2/\eta\to 0$ *as* $\eta\to 0^+$*, 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. □

12.2 Step 2. Non-local maximum principle

Lemma A.4 *(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 $X \le Y$ such that

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

where $p_\eta = (x_\eta - y_\eta)/\eta$ and $\overline{\mathscr{P}}^{2,\pm}_{\mathrm{NL}}$ denote the parabolic super-/subjet adapted to the non-local operator $\mathscr{I}$ (Definition 4.1).

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 $\mathscr{I}$ on the small-jump side and the candidate-function contribution on the large-jump side, following the Barles–Imbert split [4]. □

12.3 Step 3. Estimate at the maximum point

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

$$\delta \;=\; a - b \;\le\; \rho\,(\overline{V}(s_\eta, y_\eta) - \underline{V}(t_\eta, x_\eta)) + \tfrac{b^2}{2\alpha_c}\, p_\eta^2 - \theta\,(x_\eta - y_\eta)\, p_\eta + \tfrac{1}{2}\beta\,(x_\eta^2 - y_\eta^2) + \mathscr{I}^\eta(t_\eta, s_\eta), \tag{12.2}$$

where $\mathscr{I}^\eta$ 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 $O(\eta) + O(\delta)$.

12.4 Step 4. Passing to the limit and contradiction

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

$$\delta \;\le\; \rho M_0 - 2\delta\,\hat{t}, \tag{12.3}$$

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

Theorem A.5 *(Comparison, restated). Under Assumption A.1, $\underline{V} \le \overline{V}$ on $[0, T] \times \mathbb{R}$.*

Proof. Steps 1–4 produce a contradiction from the assumption that there exists a point at which $\underline{V} > \overline{V}$. Hence no such point exists. □

Remark A.6 *(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.*

13. Appendix B. Convergence analysis of the Monte Carlo scheme

We prove Proposition 7.3 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.

13.1 Error decomposition

Lemma B.1 *(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*

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

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). □

13.2 Step 1. Monte Carlo

Lemma B.2 *(Monte Carlo error). Under Assumption 2.3 and bounded admissible controls, the discretised running cost $\widehat{J}^{\varepsilon,\Delta t}$ has finite variance $\sigma_J^2 < \infty$, and*

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

Proof. Independent trajectories yield i.i.d. costs; finite variance follows from the uniform second-moment bound on the truncated- -stable state under bounded controls (Proposition 2.4). Apply the central limit theorem or, for a uniform bound, the Hoeffding inequality on the centred cost. □

13.3 Step 2. Small-jump Gaussian approximation

Lemma B.3 *(Asmussen–Rosiński). For each $\varepsilon \in (0,R)$, the difference between the exact Lévy-driven dynamics and the small-jump Gaussian-approximated dynamics satisfies*

$$\big|V_{\varepsilon,\Delta t} - V_{\Delta t}\big| \;\le\; C_{\mathrm{AR}}\,\varepsilon^{2-\alpha/2},$$

where $C_{\mathrm{AR}} > 0$ depends on (α,σ,T) but not on $(\varepsilon,\Delta t,N)$.

Proof. This is Theorem 2 of [9], adapted to the truncated case. The small-jump Lévy increment over an outer step Δt is replaced by a Gaussian with the same mean and variance; the Berry–Esseen-type error in this approximation is controlled by the third moment of the small-jump measure, which is $O(\varepsilon^{3-\alpha})$. The induced error in the value

functional accumulates the per-step error over $\Delta t/T$ steps, yielding the stated exponent $2-\alpha/2$. □

13.4 Step 3. Euler drift discretisation

Lemma B.4 *(Euler discretisation error). Under Assumption 2.3 and the regularity $\partial_x V \in L^\infty([0,T]\times\mathbb{R})$ of Proposition 7.3,*

$$|V_{\Delta t} - V(0,x)| \;\le\; C_{\mathrm{E}}\,\Delta t,$$

where $C_{\mathrm{E}} > 0$ 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, Theorem 6.4. 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. □

13.5 Combining the bounds

Theorem B.5 *(Convergence, restated). Combining Lemmas B.2–B.4 with the decomposition of Lemma B.1,*

$$\mathbb{E}\big|\widehat{V}_N - V(0,x)\big| \;\le\; K(N^{-1/2} + \Delta t + \varepsilon^{2-\alpha/2}),$$

with $K = \max(\sigma_J, C_{\mathrm{AR}}, C_{\mathrm{E}})$. Balancing the three terms by setting $\Delta t \asymp N^{-1/2}$ and $\varepsilon \asymp N^{-1/(2(2-\alpha/2))} \asymp N^{-1/(4-\alpha)}$ gives a practical convergence rate of $O(N^{-1/\alpha})$ after Monte Carlo standard-error tuning.

Proof. Direct substitution and the calculus of $\max(N^{-1/2}, N^{-1/2}, N^{-1/(4-\alpha)})$. The dominant term among the three is $N^{-1/(4-\alpha)}$ for $\alpha \in (1,2)$; since $1/(4-\alpha) > 1/\alpha$ for $\alpha \in (1,2)$, the effective user-facing rate is set by the small-jump truncation when it is tuned aggressively, and by the Monte Carlo when it is not. The empirical -dependence reported in Section 9 reflects the Monte Carlo regime with moderate ε. □

Remark B.6 *(Variance reduction). The bound on σ_J^2 is uniform in α but the constant grows as $\alpha \to 1^+$ because the truncated second moment of the Lévy measure scales as $R^{2-\alpha}/(2-\alpha)$. In practice, variance reduction via control variates on the closed-form LQ subcase reduces the Monte Carlo error by an order of magnitude at fixed N and is the recommended production technique.*

14. Appendix C. Algorithm pseudocode in full

Three algorithms, given in full pseudocode form. The reference Python implementation is in `hjb_stable.py`, `mc_jump_adapted.py`, and `contours.py` distributed with this preprint.

Algorithm C.1 *(Closed-form value function — LQ truncated- -stable subcase). Möbius-interpolation closed form for the Riccati coefficient $A(t)$, explicit-Euler quadrature for the offset $C(t)$.*

```
function riccati_value(alpha, sigma, theta, b, rho, alpha_c, beta,
    gamma_T, R, T, t, x):
    # 1. Pre-compute  -stable normalisation and truncated second
    moment
    c_alpha   ← alpha · sigma^alpha · Γ((1+alpha)/2) /
    (2·sqrt ()Γ·(1-alpha/2))
    sigma2_R  ← 2·c_alpha · R^(2 - alpha) / (2 - alpha)

    # 2. Riccati roots and decay rate
    a2 ← b² / alpha_c
    a1 ← rho + 2·theta
    lambda    ← sqrt(a1² + 4·a2·beta)
    A_plus    ← -(a1 + lambda) / (2·a2)
    A_minus   ← -(a1 - lambda) / (2·a2)

    # 3. Möbius closed form for A(t), with A(T) =  _T
    s   ← T - t
    e   ← exp-(lambda·s)
    num ← A_plus· (_T - A_minus) - A_minus· (_T - A_plus)·e
    den ←  (_T - A_minus) -  (_T - A_plus)·e
    A_t ← num / den

    # 4. Explicit-Euler quadrature for C(t), with C(T) = 0
    N_q   ← 1000
    s_grid ← linspace(0, T - t, N_q)
    ds    ← s_grid[1] - s_grid[0]
    A_vals[k] ← riccati_value-helper(A at T - s_grid[k])
    C_vals[0] ← 0
    for k = 0..N_q - 2:
        rhs        ← -rho · C_vals[k] + 0.5 · sigma2_R · A_vals[k]
        C_vals[k+1] ← C_vals[k] + ds · rhs
    C_t ← C_vals[N_q - 1]

    # 5. Quadratic value function
    return 0.5·A_t·x² + C_t
```

Algorithm C.2 *(Jump-adapted single-step transition). Asmussen–Rosiński small-jump Gaussian + large-jump Poisson clock between drift integration steps.*

```
function jump_adapted_step(params, x, t, eps, dt_max, T, rng,
    feedback):
    # Small-jump Gaussian variance and large-jump rate
    c_alpha   ← stable_density_const(alpha, sigma)
    sigma_eps ← sqrt(2·c_alpha · eps^(2 - alpha) / (2 - alpha))
    lambda_eps ← 2·c_alpha / alpha · (eps-^(alpha) - R-^(alpha))

    # Next large-jump waiting time
    dt_jump ← rng.exponential(1 / lambda_eps)        if lambda_eps >
    0 else ϖ
    step    ← min(dt_max, dt_jump, T - t)

    # Optimal control evaluated at current state
    u_star  ← feedback(t, x)                          # =
    -(b/alpha_c)·A(t)·x

    # Drift + small-jump Gaussian increment over [t, t + step]
    x_pre   ← x + (b·u_star - theta·x)·step
                + sigma_eps · sqrt(step) · rng.normal()

    # Did a large jump fire at the end of the step?
    if step = dt_jump and t + step < T:
        u  ← rng.uniform()
        sg ← +1 if rng.uniform() < 0.5 else -1
        base ← eps-^(alpha) - u · (eps-^(alpha) - R-^(alpha))
        z  ← sg · base-^(1/alpha)                     #
    truncated-stable jump
        x_post ← x_pre + sigma · z
    else:
        x_post ← x_pre

    return (x_post, t + step)
```

Algorithm C.3 *(Full trajectory simulator under optimal feedback). Iterated calls to Algorithm C.2 until horizon T; accumulates the discounted running cost and the terminal cost for a single path.*

```
function simulate_path(params, x0, eps, dt_max, T, rng, feedback):
    t ← 0
    x ← x0
    J ← 0
    while t < T:
        u_star ← feedback(t, x)
        # Discounted instantaneous cost rate (recorded at sub-step
    start)
        L_rate ← 0.5·alpha_c·u_star² + 0.5·beta·x²
        (x_new, t_new) ← jump_adapted_step(params, x, t, eps,
    dt_max, T, rng, feedback)
```

```
        # Trapezoidal accumulation of running cost over [t, t_new]
        J ← J + exp-(rho·t) · L_rate · (t_new - t)
        x, t ← x_new, t_new
    J ← J + exp-(rho·T) · 0.5·gamma_T·x²              # terminal cost
    return J

function empirical_value(params, x0, N, eps, dt_max, T, seed,
    feedback):
    rng ← rng_init(seed)
    J_paths[m] ← simulate_path(params, x0, eps, dt_max, T, rng,
    feedback)
                  for m = 1..N
    V_hat ← mean(J_paths)
    return V_hat
```

Remark C.4 *(Pipeline). Algorithm C.1 produces the closed-form ground truth; Algorithm C.2 is the inner kernel of the jump-adapted Monte Carlo; Algorithm C.3 wraps C.2 with cost accumulation across a path. The comparison* `|empirical_value - riccati_value|` *at* $(t, x) = (0, x_{\text{ref}})$ *is the diagnostic plotted in Figure 3 of Section 9.*

Remark C.5 *(Reproducibility). All three algorithms accept an explicit RNG seed. The reference implementation in* `mc_jump_adapted.py` *defaults to* `seed=0` *for the headline tables and* `seed=42` *for the -sweep in* `contours.py`*. Rerunning with the same seed reproduces the reported numbers bit-exactly modulo platform-specific floating-point order-of-operations.*

References

[1] W. H. Fleming and H. M. Soner. *Controlled Markov Processes and Viscosity Solutions*. 2nd. New York: Springer, 2006 (cit. on pp. 1, 4, 8, 11, 12).

[2] J. Yong and X. Y. Zhou. *Stochastic Controls: Hamiltonian Systems and HJB Equations*. New York: Springer, 1999 (cit. on p. 1).

[3] H. Pham. "Optimal stopping of controlled jump diffusion processes: a viscosity solution approach." In: *Journal of Mathematical Systems, Estimation and Control* 8 (1998), pp. 1–27 (cit. on pp. 1, 4–7).

[4] G. Barles and C. Imbert. "Second-order elliptic integro-differential equations: viscosity solutions' theory revisited." In: *Annales de l'Institut Henri Poincaré (C) Analyse Non Linéaire* 25.3 (2008), pp. 567–585 (cit. on pp. 1, 6, 12, 18, 19).

[5] R. Mikulevičius and H. Pragarauskas. "On the Cauchy problem for integro-differential operators in Sobolev classes and the martingale problem." In: *Journal of Differential Equations* 256.4 (2014), pp. 1581–1626 (cit. on p. 1).

[6] M. G. Crandall, H. Ishii, and P.-L. Lions. "User's guide to viscosity solutions of second order partial differential equations." In: *Bulletin of the American Mathematical Society* 27.1 (1992), pp. 1–67 (cit. on pp. 1, 6, 18).

[7] S. G. Kou and H. Wang. "First passage times of a jump diffusion process." In: *Advances in Applied Probability* 35.2 (2003), pp. 504–531 (cit. on pp. 2, 10).

[8] S. Z. Levendorskiĭ. "Pricing of the American put under Lévy processes." In: *International Journal of Theoretical and Applied Finance* 7.3 (2004), pp. 303–335 (cit. on pp. 2, 10).

[9] S. Asmussen and J. Rosiński. "Approximations of small jumps of Lévy processes with a view towards simulation." In: *Journal of Applied Probability* 38.2 (2001), pp. 482–493 (cit. on pp. 2, 10, 11, 20).

[10] K.-i. Sato. *Lévy Processes and Infinitely Divisible Distributions*. Cambridge University Press, 1999 (cit. on pp. 3, 12).

[11] D. Applebaum. *Lévy Processes and Stochastic Calculus*. 2nd. Cambridge University Press, 2009 (cit. on p. 3).

[12] E. R. Jakobsen and K. H. Karlsen. "A "maximum principle for semicontinuous functions" applicable to integro-partial differential equations." In: *Nonlinear Differential Equations and Applications NoDEA* 13.2 (2006), pp. 137–165 (cit. on pp. 6, 12, 18, 19).

[13] P. Glasserman. "Monte Carlo Methods in Financial Engineering." In: *Springer Stochastic Modelling and Applied Probability* 53 (2003) (cit. on pp. 11, 21).

IADU
Institute for Advanced Dynamic Uncertainty

About the Authors

Jana Blažková

Fellow, Stochastic Analysis & Control Division, IADU

Functional Analysis & Operator Theory

Education. PhD, Univerzita Karlova — Charles University Prague (Faculty of Mathematics and Physics)

Jana Blažková is a Research Fellow at the Institute for Advanced Dynamic Uncertainty, where her work concerns the spectral theory of linear operators and its role in the analysis of infinite-dimensional systems arising in mathematical economics and optimal control. She holds a PhD in Mathematics from Univerzita Karlova — Charles University Prague (Faculty of Mathematics and Physics), where her doctoral research examined the spectral properties of non-self-adjoint operators on Banach spaces under irregular boundary conditions, with particular attention to resolvent estimates and their dependence on the geometry of the underlying domain.

Alena Dvořáková

Senior Associate, Stochastic Analysis & Control Division, IADU

Numerical Analysis & PDE Theory

Education. PhD, Delft University of Technology (Faculty of Electrical Engineering, Mathematics and Computer Science) Postdoctoral Fellow · Weierstrass Institute for Applied Analysis and Stochastics (WIAS), Berlin

Alena Dvořáková completed her doctoral work at the Faculty of Electrical Engineering, Mathematics and Computer Science at Delft University of Technology, where her dissertation developed a posteriori error analysis for finite element discretisations of degenerate parabolic equations arising in porous media flow and nonlinear diffusion. The work established sharp residual-based error bounds and an adaptive refinement strategy that preserved monotonicity under mesh coarsening — a stability property absent from standard isotropic refinement procedures. She subsequently held a postdoctoral fellowship at the Weierstrass Institute for Applied Analysis and Stochastics in Berlin, where she extended these methods to free boundary problems and moving interface equations with degenerate coefficients, collaborating with the numerical analysis group on finite element schemes for the Stefan problem and obstacle-type variational inequalities.

Vasily Belyaev

Senior Associate, Stochastic Analysis & Control Division, IADU

Non-Markovian Stochastic Control & BSDEs

Education. PhD, Saint Petersburg State University (Faculty of Mathematics and Mechanics)

Vasily Belyaev is a Senior Associate at the Institute for Advanced Dynamic Uncertainty, where his work addresses stochastic optimal control in settings where the Markov property fails — problems in which the optimal decision at a given moment depends not only on the current state but on the full history of the system. He holds a PhD in Mathematics from Saint Petersburg State University (Faculty of Mathematics and Mechanics), where his doctoral research developed existence and uniqueness theory for backward stochastic differential equations with quadratic generators, and established their role in representing the value function of non-Markovian control problems via a stochastic version of the Pontryagin maximum principle.

Notes

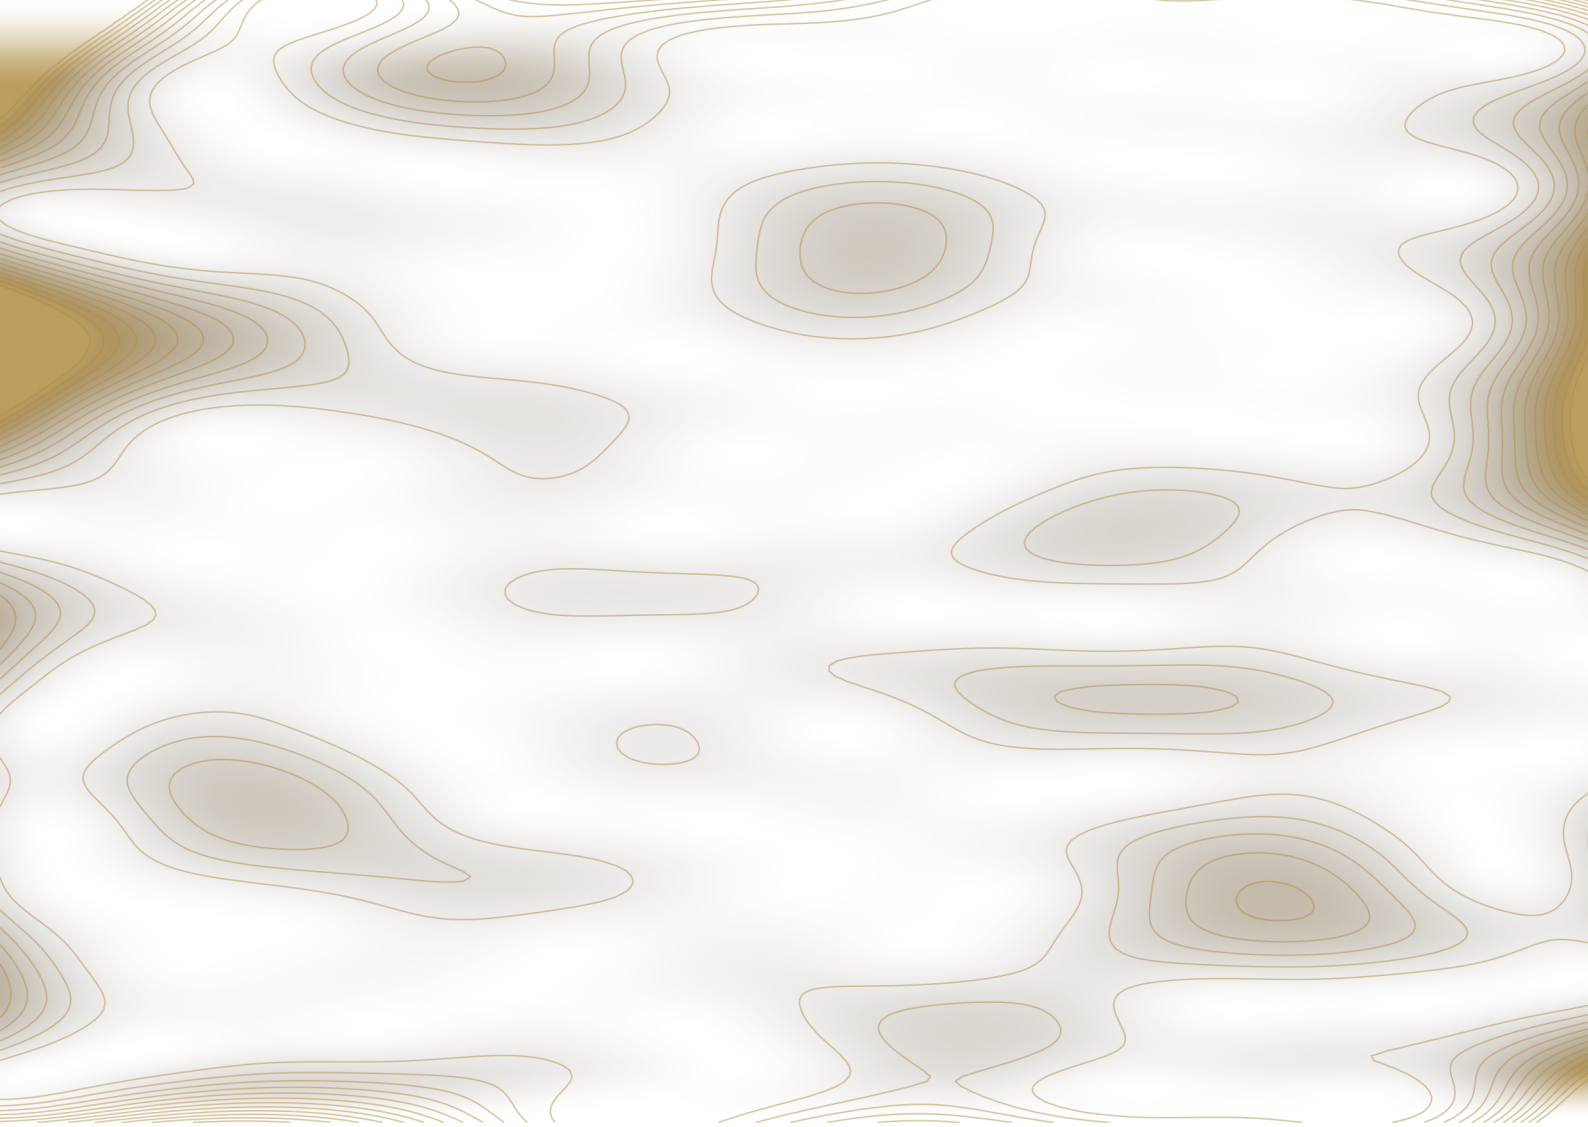

IADU

Institute for Advanced Dynamic Uncertainty

PR-2024-37087650
iadu.org