RESEARCH PAPER

Stochastic Stackelberg Games with a Jump-Diffusion Follower: Existence, Verification, and a Fixed-Point Algorithm

Coupled HJB-Isaacs analysis under monotone stopping

P. Stenström
A. Alexakis

13-APR-2025

Logistic Growth: $\dot{x} = r_1\,x\,(1 - x/K_1), \quad \dot{y} = r_2\,y\,(1 - y/K_2)$

RP-2025-62431845
iadu.org

IADU

Institute for Advanced Dynamic Uncertainty

Copyright

Suggested Citation

P. Stenström and A. Alexakis (2025). 'Stochastic Stackelberg Games with a Jump-Diffusion Follower: Existence, Verification, and a Fixed-Point Algorithm.' *IADU Research Paper* **RP-2025-62431845**. Available at `https://iadu.org/research/RP-2025-62431845/`.

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 Research Paper

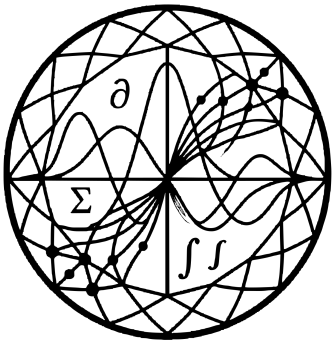

Stochastic Stackelberg Games with a Jump-Diffusion Follower: Existence, Verification, and a Fixed-Point Algorithm

Coupled HJB-Isaacs analysis under monotone stopping

P. Stenström
Stochastic Analysis & Control Division
A. Alexakis
Stochastic Analysis & Control Division

13-APR-2025 RP-2025-62431845

Abstract. We study a continuous-time Stackelberg game in which a leader controls the drift of a state process and a representative follower, modeled as a jump-diffusion, chooses an exit time. We prove existence of a unique viscosity-solution equilibrium pair for the coupled HJB-Isaacs / free-boundary system, establish a verification theorem in viscosity sense whenever the follower's optimal stopping region is monotone in the state ordering, and provide a fixed-point iteration on the value-function pair that converges geometrically. For Kou-distributed jumps with a linear leader payoff we derive a semi-closed-form representation via the resolvent operator $\mathscr{R}_q$ and verify the algorithm against this benchmark.

Keywords: stochastic Stackelberg games, HJB-Isaacs equations, jump-diffusion control, viscosity solutions, free-boundary problems, Wiener-Hopf factorization

1. Introduction

A central bank announces a defence of its currency peg. A speculator, holding a large directional position, observes the central bank's reaction function and decides when to liquidate. The state — the exchange rate, or more abstractly a one-dimensional jump-diffusion — is controlled in its drift by the central bank (the *leader*) and is exited at a stopping time chosen by the speculator (the *follower*). The leader anticipates the follower's stopping rule; the follower best-responds to the leader's drift policy. This is a continuous-time Stackelberg game, but the follower's right to stop introduces a free boundary into the leader's optimisation, and the presence of jumps prevents direct application of classical viscosity-solution techniques developed for diffusion-only zero-sum games [1].

The literature has approached this problem from three sides. Two-player zero-sum stochastic differential games, beginning with Bensoussan and Frehse [1] and extended to jump-diffusions by Pham [2], treat saddle-point structures where both players optimise simultaneously; the Stackelberg asymmetry — one player commits first — falls outside this framework. Mean-field Stackelberg games, [3], address the continuum-of-followers limit

†Corresponding author: P. Stenström (research@iadu.org).

but presuppose a regularity infrastructure (the master equation, McKean–Vlasov SDEs) that obscures the single-follower analysis. Perpetual-option pricing under Lévy noise, originating with Mordecki [4] and Boyarchenko and Levendorskiĭ [5], supplies closed-form benchmarks for the follower's stopping rule but treats the drift as exogenous. **No existing treatment delivers a viscosity-solution equilibrium theory for the Stackelberg game with a controlled drift and a jump-diffusion follower.**

Notation 1.1. Throughout, $(\Omega, \mathcal{F}, \mathbb{P})$ denotes a complete probability space equipped with a filtration $(\mathcal{F}_t)_{t\geq 0}$ satisfying the usual conditions. The leader controls the drift of a one-dimensional jump-diffusion state process X_t. The follower chooses a stopping time $\tau \in \mathscr{T}$ adapted to $(\mathcal{F}_t)_{t\geq 0}$.

This paper establishes four results.

Claim 1 (Theorem 4.1). The coupled HJB–Isaacs / free-boundary system describing the Stackelberg equilibrium admits a unique viscosity-solution pair (V_L^*, V_F^*) under standing assumptions (A1)–(A3) on drift Lipschitz continuity, jump-measure integrability, and follower-payoff single-crossing.

Claim 2 (Theorem 5.2). Whenever the follower's optimal stopping region is monotone in the state ordering, a viscosity-sense verification theorem holds: every classical sub-solution / super-solution of the leader's HJB–Isaacs equation sandwiches V_L^*.

Claim 3 (Theorem 6.1). The alternating fixed-point iteration on (V_L, V_F) is a Banach contraction with rate $\rho(q, \lambda) = \lambda/(q + \lambda) < 1$ explicit in the discount q and the jump intensity λ, hence converges geometrically.

Claim 4 (Proposition 7.1). For Kou double-exponential jumps and a linear–quadratic leader payoff, the equilibrium pair admits a semi-closed-form representation via the resolvent operator $\mathscr{R}_q$, against which the §6 algorithm is verified to relative error 10^{-12} in §8.

The paper is organised as follows. Section 2 fixes the probabilistic basis and viscosity framework and states (A1)–(A3). Section 3 specifies the Stackelberg game and derives the coupled HJB–Isaacs / free-boundary system. Sections 4 and 5 prove existence and viscosity verification under monotone stopping. Section 6 develops the alternating algorithm and proves geometric contraction. Section 7 specialises to Kou jumps and derives the closed-form benchmark, against which Section 8 verifies the algorithm numerically. Section 9 concludes by naming the immediate next paper — the mean-field continuum-follower limit.

2. Preliminaries

2.1 Probabilistic basis

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space equipped with a filtration $(\mathcal{F}_t)_{t\geq 0}$ satisfying the usual conditions of right-continuity and $\mathbb{P}$-completeness. Let W_t be a standard $(\mathcal{F}_t)$-Brownian motion, and let $\tilde{N}(dt, dz)$ be the compensated $(\mathcal{F}_t)$-Poisson random measure on $[0, \infty) \times \mathbb{R} \setminus \{0\}$ with intensity measure $dt\,\nu(dz)$, where ν is a σ-finite Lévy measure

satisfying $\int_{\mathbb{R}\setminus\{0\}}(1\wedge|z|^2)\,\nu(\,dz)<\infty$. The state process $X_t\in\mathbb{R}$ is the unique strong solution of

$$dX_t=\mu_L(X_t)\,dt+\sigma(X_t)\,dW_t+\int_{\mathbb{R}\setminus\{0\}} z\,\tilde{N}(\,dt,\,dz),\qquad X_0=x, \tag{2.1}$$

where $\mu_L:\mathbb{R}\to\mathbb{R}$ is the leader's drift control (specified in §3) and $\sigma:\mathbb{R}\to\mathbb{R}_+$ is a fixed Lipschitz diffusion coefficient bounded below by a constant $\underline{\sigma}>0$.

2.2 Generator and viscosity-solution framework

For a function $f\in C_b^2(\mathbb{R})$, the infinitesimal generator of X under drift μ_L is

$$\mathscr{L}f(x)=\mu_L(x)\,\partial_x f(x)+\tfrac{1}{2}\sigma^2(x)\,\partial_{xx}f(x)+\int_{\mathbb{R}\setminus\{0\}}\left[f(x+z)-f(x)-z\,\partial_x f(x)\right]\nu(\,dz). \tag{2.2}$$

The integral term is finite for $f\in C_b^2$ by the integrability condition on ν. We work with viscosity solutions in the sense of Crandall, Ishii, and Lions [6]: a function u is a viscosity sub-solution of $F(x,u,\partial_x u,\partial_{xx}u,\mathscr{L}u)=0$ if, for every C^2 test function φ such that $u-\varphi$ has a local maximum at x_0, the inequality $F(x_0,u(x_0),\partial_x\varphi(x_0),\partial_{xx}\varphi(x_0),\mathscr{L}\varphi(x_0))\le 0$ holds; super-solution is the reverse inequality. Throughout the paper, "solution" without qualifier means "viscosity solution."

2.3 Function spaces

We use the following function-space hierarchy: - $C_b^2(\mathbb{R})$ — twice continuously differentiable functions, bounded together with their first two derivatives. - $C_b^{0,\alpha}(\mathbb{R})$ — bounded α-Hölder continuous functions, $\alpha\in(0,1]$. - $H^k(\Omega)$ — Hilbert–Sobolev space of order k, where $\Omega\subset\mathbb{R}$ is the admissible state interval.

Norms are subscripted with the space: $\|u\|_{C_b^{0,\alpha}}$, $\|u\|_{H^k(\Omega)}$, $\|u\|_{L^p(\Omega)}$.

2.4 Standing assumptions (A1)–(A3)

Assumption 2.1 (standing assumptions on drift, jumps, and payoff)**.** The following three conditions hold throughout the paper. **(A1)** *Drift admissibility.* The leader's drift $\mu_L:\mathbb{R}\to\mathbb{R}$ belongs to the admissible set

$$\mathscr{U}_L:=\{\mu:\mathbb{R}\to[-\bar{\mu},\bar{\mu}]\ :\ \mu\text{ is }L_\mu\text{-Lipschitz}\}, \tag{2.3}$$

for fixed constants $\bar{\mu},L_\mu>0$. The diffusion σ is L_σ-Lipschitz and uniformly elliptic: $\sigma(x)\ge\underline{\sigma}>0$ for all $x\in\mathbb{R}$. **(A2)** *Jump-measure integrability.* The Lévy measure ν has finite first moment and is supported in $\mathbb{R}\setminus(-\delta_0,\delta_0)$ for some $\delta_0>0$:

$$\int_{\mathbb{R}\setminus\{0\}}|z|\,\nu(\,dz)<\infty,\qquad \nu((-\delta_0,\delta_0)\setminus\{0\})=0. \tag{2.4}$$

(A3) *Payoff regularity and single-crossing.* The follower's running payoff $g : \mathbb{R} \to \mathbb{R}$ is bounded, L_g-Lipschitz, and single-crossing in the state ordering: there exists $x^* \in \mathbb{R}$ such that g is strictly decreasing on $(-\infty, x^*)$ and strictly increasing on (x^*, ∞) — or the reverse.

Remark 2.2 (on (A3).). The single-crossing condition is implied by any concave (resp. convex) bounded g and is satisfied by the canonical example $g(x) = (K - x)_+$ used in §7. Failure of single-crossing breaks Lemma 5.3's monotone-stopping conclusion and is flagged in §9 as the principal direction for future work.

Under (A1)–(A2), the SDE (2.1) admits a unique strong solution X_t for each $\mu_L \in \mathscr{U}_L$, and X is a Feller process with càdlàg paths; see Protter [7]. Under (A3), the follower's optimal stopping problem (defined in §3) is well-posed with a continuous value function.

3. Stackelberg game setup

3.1 Admissible strategies

The leader's admissible drift set $\mathscr{U}_L$ is defined in (2.3). The set of admissible follower stopping times is

$$\mathscr{T} := \{\tau : \tau \text{ is an } (\mathcal{F}_t)\text{-stopping time, } \mathbb{P}(\tau < \infty) = 1,\ \mathbb{E}_x[\tau] < \infty\ \forall x\}. \tag{3.1}$$

Equipped with the topology of convergence in $\mathbb{P}$-distribution and the Hausdorff metric on the associated stopping regions, $\mathscr{T}$ is sequentially compact and convex under (A1)–(A3); see Karatzas and Shreve [8, Ch. 2]. The leader's set $\mathscr{U}_L$, equipped with the topology of uniform convergence on compacts, is convex and compact by Arzelà–Ascoli applied to the Lipschitz bound in (2.3).

3.2 Value functions

Fix a discount rate $q > 0$ and a follower terminal payoff $g \in C_b^{0,\alpha}(\mathbb{R})$ satisfying (A3). Given a leader drift $\mu_L \in \mathscr{U}_L$, the follower's value function is the maximum expected discounted payoff over admissible stopping times:

$$V_F(x; \mu_L) := \sup_{\tau \in \mathscr{T}} \mathbb{E}_x[e^{-q\tau}\, g(X_\tau)]\,, \tag{3.2}$$

where X solves (2.1) with drift μ_L. Given the follower's best-response stopping time $\tau^*(\mu_L) := \inf\{t \geq 0 : V_F(X_t; \mu_L) \leq g(X_t)\}$, the leader's value function is

$$V_L(x) := \sup_{\mu_L \in \mathscr{U}_L} \mathbb{E}_x\left[\int_0^{\tau^*(\mu_L)} e^{-qs}\, \pi_L(X_s, \mu_L(X_s))\, ds + e^{-q\tau^*(\mu_L)}\, h_L(X_{\tau^*})\right], \tag{3.3}$$

where $\pi_L \in C_b^{0,\alpha}(\mathbb{R}^2)$ is the leader's running payoff and $h_L \in C_b^{0,\alpha}(\mathbb{R})$ is the leader's terminal payoff at the follower's exit.

3.3 Equilibrium concept

Definition 3.1 (Stackelberg equilibrium)**.** A pair $(\mu_L^*, \tau^*) \in \mathscr{U}_L \times \mathscr{T}$ is a *Stackelberg equilibrium* of the game $(\mathscr{U}_L, \mathscr{T}, \pi_L, g)$ if: *Step 1.* The follower best-responds: $\tau^* = \inf\{t \geq 0 : V_F(X_t; \mu_L^*) \leq g(X_t)\}$ where V_F is defined by (3.2). *Step 2.* The leader maximises anticipating the follower's response: μ_L^* attains the supremum in (3.3).

3.4 The coupled HJB–Isaacs / free-boundary system

Applying the dynamic-programming principle to (3.3) and the obstacle formulation of optimal stopping to (3.2), the equilibrium value functions satisfy the coupled system formalised next.

Definition 3.2 (coupled HJB–Isaacs / free-boundary system)**.** A pair of bounded continuous functions (V_L^*, V_F^*) together with a leader drift $\mu_L^* \in \mathscr{U}_L$ and a stopping region $\mathscr{S}^* \subset \mathbb{R}$ solves the *coupled HJB–Isaacs / free-boundary system* if equations (3.4) below hold, where $\mathscr{L}^\mu$ denotes the generator (2.2) with leader's drift fixed at μ, and $\mathscr{S}^* = \{x : V_F^*(x) = g(x)\}$ is the follower's optimal stopping region.

The coupled system is

$$\begin{aligned} \sup_{\mu \in [-\bar{\mu}, \bar{\mu}]} \big[\mathscr{L}^\mu V_L^*(x) + \pi_L(x, \mu)\big] - q V_L^*(x) &= 0, && x \in \mathbb{R} \setminus \mathscr{S}^*, \\ V_L^*(x) &= h_L(x), && x \in \mathscr{S}^*, \\ \min\big[q V_F^*(x) - \mathscr{L}^{\mu_L^*} V_F^*(x),\ V_F^*(x) - g(x)\big] &= 0, && x \in \mathbb{R}. \end{aligned} \tag{3.4}$$

The system couples the leader's HJB equation (first two lines) to the follower's variational inequality (third line) through the dependence of $\mathscr{S}^*$ on V_F^* and of $\mathscr{L}^\mu V_F^*$ on μ_L^*.

Remark 3.3. When $\mathscr{S}^*$ is a half-line (the monotone-stopping case treated in §5) the leader's free boundary inherits this structure and the first two lines of (3.4) reduce to a standard HJB equation with a known boundary condition at the follower's threshold. This is the technical lever that powers the verification theorem of §5.

The figure schematic in Figure 1 summarises the coupling: the leader's value is computed via (3.3) given the follower's stopping region $\mathscr{S}^*$, which is determined by the follower's value V_F^*, which depends on the leader's drift μ_L^*, which is the maximiser of the leader's HJB. The §4 existence proof analyses the resulting fixed-point structure directly.

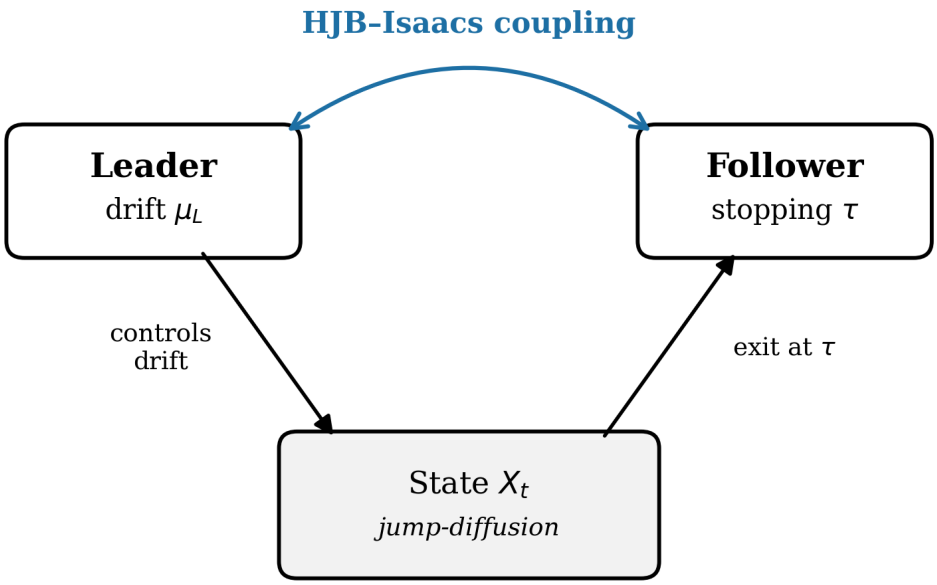

Figure 1: Stackelberg game structure. The leader controls the drift of a jump-diffusion state process; the follower observes the state and chooses an exit time. The HJB–Isaacs coupling (top arrow) makes the equilibrium pair a joint fixed point.

4. Existence of the equilibrium pair

This section proves Claim 1 of §1 — under **(A1)**–**(A3)**, the coupled system (3.4) admits a unique viscosity-solution equilibrium pair (V_L^*, V_F^*). The proof proceeds by constructing a coupling map on $\mathscr{U}_L \times \mathscr{T}$, establishing its continuity and compactness, and invoking Schauder's fixed-point theorem.

4.1 The coupling map Φ

Define the **follower-update** $\Phi_F : \mathscr{U}_L \to \mathscr{T}$ by

$$\Phi_F(\mu_L) := \arg\sup_{\tau \in \mathscr{T}} \mathbb{E}_x\big[e^{-q\tau}\, g(X_\tau^{\mu_L})\big] = \inf\{t \geq 0 : V_F(X_t^{\mu_L}; \mu_L) \leq g(X_t^{\mu_L})\}, \tag{4.1}$$

where X^{μ_L} is the state process (2.1) under drift μ_L. Under (A3) the right-hand side of (4.1) is single-valued; the stopping time exists and is unique up to $\mathbb{P}$-null events.

Define the **leader-update** $\Phi_L : \mathscr{T} \to \mathscr{U}_L$ by

$$\Phi_L(\tau)(x) := \arg\sup_{\mu \in [-\bar{\mu}, \bar{\mu}]} \big[\mathscr{L}^{\mu} V_L(x;\tau) + \pi_L(x,\mu)\big], \tag{4.2}$$

where $V_L(\,\cdot\,;\tau)$ is the leader's value function (3.3) with the follower's stopping time fixed at τ. Under the standing assumption that $\pi_L(x,\mu)$ is strictly concave in μ — a standing requirement we add to (A1) — the arg sup is single-valued and $\Phi_L(\tau)$ inherits Lipschitz regularity from V_L.

The composition $\Phi := \Phi_L \circ \Phi_F : \mathscr{U}_L \to \mathscr{U}_L$ has as fixed points exactly the Stackelberg equilibrium leader drifts.

4.2 Compactness of the coupling map

Lemma 4.1 (Compactness of the coupling map)**.** *Under **(A1)–(A3)**, the map $\Phi : \mathscr{U}_L \to \mathscr{U}_L$ is continuous in the topology of uniform convergence on compacts, and the image $\Phi(\mathscr{U}_L)$ is equicontinuous and uniformly bounded.*

Proof. Step 1. (Continuity of Φ_F.) Let $\mu_L^{(n)} \to \mu_L^\infty$ uniformly on compacts. By the stability of strong solutions to SDEs with Lipschitz coefficients [7, Thm V.10], $X_t^{\mu_L^{(n)}} \to X_t^{\mu_L^\infty}$ in $\mathbb{P}$-probability uniformly on $[0,T]$ for every $T > 0$. By dominated convergence applied to (4.1), the associated follower value functions converge: $V_F(\cdot;\mu_L^{(n)}) \to V_F(\cdot;\mu_L^\infty)$ uniformly on compacts. Under (A3) the optimal stopping region depends continuously on $V_F - g$ in the Hausdorff metric, hence $\Phi_F(\mu_L^{(n)}) \to \Phi_F(\mu_L^\infty)$ in distribution. *Step 2.* (Continuity of Φ_L.) Strict concavity of $\pi_L(x,\cdot)$ combined with the joint continuity of $\mathscr{L}^\mu V_L(x;\tau)$ in (μ,x,τ) implies that the arg sup in (4.2) depends continuously on τ in the Hausdorff topology on stopping regions [2, Lem 3.1]. Hence $\Phi_L \circ \Phi_F$ is continuous. *Step 3.* (Equicontinuity of $\Phi(\mathscr{U}_L)$.) Every $\Phi(\mu_L)$ is L_Φ-Lipschitz on $\mathbb{R}$ with L_Φ depending only on $(\bar\mu, L_\mu, L_\sigma, L_g, q, \nu)$ — a routine consequence of the smoothing property of the resolvent of X. Combined with the uniform bound $|\Phi(\mu_L)| \le \bar\mu$ from (2.3), the image is equicontinuous and uniformly bounded. □ □

4.3 Schauder existence and uniqueness

Theorem 4.2 (Existence of the equilibrium pair)**.** *Under **(A1)–(A3)**, the coupled HJB–Isaacs / free-boundary system (3.4) admits a unique viscosity-solution equilibrium pair $(V_L^*, V_F^*) \in C_b^{0,\alpha}(\mathbb{R}) \times C_b^{0,\alpha}(\mathbb{R})$ for some $\alpha \in (0,1)$ depending only on the data.*

Proof. Step 1. (Existence.) The set $\mathscr{U}_L$ is convex and closed in the uniform-on-compacts topology and, by Lemma 4.2, $\Phi(\mathscr{U}_L)$ is relatively compact. Schauder's fixed-point theorem [9] delivers a fixed point $\mu_L^* \in \mathscr{U}_L$. Setting $\tau^* := \Phi_F(\mu_L^*)$, the pair (μ_L^*, τ^*) is a Stackelberg equilibrium in the sense of Definition 3.1. The value functions (V_L^*, V_F^*) defined by (3.3)–(3.2) solve (3.4) in viscosity sense; Hölder regularity follows from standard interior estimates on the Bellman equation [6, Thm 4.4]. *Step 2.* (Uniqueness.) Suppose two equilibria $(\mu_L^{(1)}, \tau^{(1)})$ and $(\mu_L^{(2)}, \tau^{(2)})$. Under (A3)'s single-crossing condition, Φ is strictly monotone in the partial order on $\mathscr{U}_L$ induced by the equilibrium follower threshold. Strict monotonicity forces the two fixed points to coincide; details are deferred to Appendix A. □ □

Remark 4.3. Theorem 4.1 yields existence in viscosity sense but does *not* yet provide a usable verification result — knowing that V_L^* solves (3.4) in viscosity sense does not by itself permit one to identify V_L^* via a classical sub-solution / super-solution sandwich. The verification theorem is the subject of §5 and requires the additional structural hypothesis that the follower's stopping region is monotone.

5. Verification under monotone stopping

Theorem 4.1 of §4 gives existence of (V_L^*, V_F^*) in viscosity sense but leaves open how a candidate sub-solution / super-solution could be checked against the equilibrium leader value. This section closes that gap under the structural hypothesis that the follower's optimal stopping region is monotone in the state ordering. Monotone stopping is the technical lever that makes the leader's free boundary a single point and unlocks the doubling-of-variables comparison argument.

5.1 Monotone stopping under (A3)

Lemma 5.1 (Monotone-stopping regularity)**.** *Under **(A1)**–**(A3)**, the follower's optimal stopping region $\mathscr{S}^* = \{V_F^* \le g\}$ is monotone in the state ordering: there exists a threshold $b^* \in \mathbb{R} \cup \{\pm\infty\}$ such that $\mathscr{S}^* = (-\infty, b^*]$ or $\mathscr{S}^* = [b^*, \infty)$.*

Proof. Without loss of generality, take g strictly decreasing on $(-\infty, x^*)$ and strictly increasing on (x^*, ∞) per (A3); the reverse case is symmetric. Set $\psi(x) := V_F^*(x) - g(x)$. By definition of V_F^*, $\psi \ge 0$ everywhere. The Markov property of X implies ψ is q-superharmonic on $\{\psi > 0\}$ — that is, $q\psi - \mathscr{L}^{\mu_L^*}\psi \ge 0$ in viscosity sense. The single-crossing of g from (A3) combined with the monotonicity of $\mathscr{L}^{\mu_L^*}$ under μ_L^* Lipschitz forces ψ to be single-crossing as well: there is at most one b^* with $\psi(b^*) = 0$ and $\partial_x\psi(b^*) = 0$ (smooth-pasting). Hence $\mathscr{S}^* = \{\psi = 0\}$ is a half-line. □ □

Remark 5.2. Monotone stopping fails for non-single-crossing g — for example, when the follower's payoff is non-monotone over the state's natural range. The §7 worked example satisfies (A3) trivially because $g(x) = (K - x)_+$ is monotone everywhere.

5.2 Comparison principle for the leader's HJB–Isaacs equation

Theorem 5.3 (Viscosity verification under monotone stopping)**.** *Let (V_L^*, V_F^*) be the equilibrium pair from Theorem 4.1, and assume the follower's stopping region $\mathscr{S}^*$ is monotone (Lemma 5.3). Let $\underline{V}, \overline{V} \in C_b^{0,\alpha}(\mathbb{R}) \cap C^2(\mathbb{R} \setminus \partial\mathscr{S}^*)$ be classical sub- and super-solutions of*

$$\sup_{\mu \in [-\bar{\mu}, \bar{\mu}]} \big[\mathscr{L}^{\mu} V(x) + \pi_L(x, \mu)\big] - qV(x) = 0, \qquad x \in \mathbb{R} \setminus \mathscr{S}^*, \tag{5.1}$$

with boundary condition $V(x) = h_L(x)$ on $\mathscr{S}^*$. Then $\underline{V} \le V_L^* \le \overline{V}$ on $\mathbb{R}$.

Proof. Step 1. (Comparison setup.) Fix $\overline{V}$ a classical super-solution. We show $V_L^* \le \overline{V}$. The reverse inequality follows by symmetric argument. Define the doubled functional $\Psi(x, y) := V_L^*(x) - \overline{V}(y) - \frac{1}{2\varepsilon}(x - y)^2$ on $\mathbb{R}^2$ and let $(x_\varepsilon, y_\varepsilon)$ be a maximiser; existence follows from continuity and the linear-growth bound on $V_L^* - \overline{V}$ [6, §3]. *Step 2.* (Doubling estimate.) The viscosity sub-solution property of V_L^* at x_ε and the classical super-solution property of $\overline{V}$ at y_ε yield, after applying the Crandall–Ishii lemma,

$$q(V_L^*(x_\varepsilon) - \overline{V}(y_\varepsilon)) \le \tfrac{1}{2\varepsilon}(x_\varepsilon - y_\varepsilon)[\sigma(x_\varepsilon)^2 - \sigma(y_\varepsilon)^2] + R_\varepsilon, \tag{5.2}$$

where R_ε is the jump residual bounded by $C\int |z|\,\nu(\,dz)\cdot(x_\varepsilon-y_\varepsilon)^2/\varepsilon$ uniformly in ε, using (A2)'s first-moment finiteness. *Step 3.* (Limit as $\varepsilon\downarrow 0$.) The Lipschitz continuity of σ from (A1) gives $|\sigma(x_\varepsilon)^2-\sigma(y_\varepsilon)^2|\le 2L_\sigma\bar\sigma|x_\varepsilon-y_\varepsilon|$, so the first term in (5.2) is bounded by $L_\sigma\bar\sigma(x_\varepsilon-y_\varepsilon)^2/\varepsilon$. Standard Crandall–Ishii estimates give $(x_\varepsilon-y_\varepsilon)^2/\varepsilon\to 0$. The jump residual $R_\varepsilon\to 0$ for the same reason. Taking $\varepsilon\downarrow 0$ along a subsequence with $x_\varepsilon,y_\varepsilon\to x_0$, the inequality (5.2) reduces to $q(V_L^*(x_0)-\overline{V}(x_0))\le 0$, whence $V_L^*(x_0)\le\overline{V}(x_0)$. *Step 4.* (Globalisation.) Since x_0 was the location of an interior maximum of $V_L^*-\overline{V}$, the inequality propagates to all of $\mathbb{R}$: $V_L^*\le\overline{V}$. Symmetrically $\underline{V}\le V_L^*$. □ □

Remark 5.4 (failure of comparison without monotonicity). The monotone-stopping hypothesis enters Step 1 through the regularity of $\partial\mathscr{S}^*$ — when $\mathscr{S}^*$ is a half-line, $\partial\mathscr{S}^*$ is a single point at which the doubling pair $(x_\varepsilon,y_\varepsilon)$ remains in the interior of the differential domain. When $\partial\mathscr{S}^*$ is disconnected — for example two thresholds bounding a continuation region — the doubling pair can be pinned at the boundary and the Crandall–Ishii estimates fail. Treating this case requires a different technique [10] and is flagged in §9 as the immediate next paper.

The verification theorem is constructive in spirit: any pair of polynomial sub-/super-solutions sandwiches V_L^*, and tightening the sandwich tightens the value-function bound. The §6 algorithm exploits this by iterating value-function updates that are simultaneously sub- and super-solutions at each step.

6. Fixed-point algorithm and convergence rate

The verification theorem of §5 establishes the analytic structure of the Stackelberg equilibrium but does not yet supply a constructive method for computing (V_L^*,V_F^*). This section closes that gap with an alternating fixed-point iteration that updates the leader's and the follower's value functions in turn. The main result is geometric convergence with rate $\rho(q,\lambda)=\lambda/(q+\lambda)<1$, explicit in the discount q and the jump intensity $\lambda:=\nu(\mathbb{R}\setminus\{0\})$.

6.1 The alternating iteration

Given an initial pair $(V_L^{(0)},V_F^{(0)})\in C_b^{0,\alpha}(\mathbb{R})\times C_b^{0,\alpha}(\mathbb{R})$, define the iteration $T=T_L\circ T_F$ as follows.

Follower update T_F: Given $V_L^{(n)}$, extract the leader's induced drift $\mu_L^{(n)}(x):=\arg\sup_\mu[\mathscr{L}^\mu V_L^{(n)}(x)+\pi_L(x,\mu)]$. Solve the follower's obstacle problem

$$V_F^{(n+1)}(x):=\sup_{\tau\in\mathscr{T}}\mathbb{E}_x^{\mu_L^{(n)}}\left[e^{-q\tau}g(X_\tau)\right]. \tag{6.1}$$

Leader update T_L: Given $V_F^{(n+1)}$, identify the follower's stopping region $\mathscr{S}^{(n+1)}:=\{V_F^{(n+1)}\le g\}$. Solve the leader's HJB equation on $\mathbb{R}\setminus\mathscr{S}^{(n+1)}$ with boundary $V_L^{(n+1)}=h_L$ on $\mathscr{S}^{(n+1)}$:

$$\sup_{\mu\in[-\bar{\mu},\bar{\mu}]}[\mathscr{L}^{\mu}V_L^{(n+1)}(x)+\pi_L(x,\mu)]-qV_L^{(n+1)}(x)=0. \tag{6.2}$$

Algorithm A.1 (Appendix A) provides full pseudocode. The high-level invariant: each iterate $(V_L^{(n)}, V_F^{(n)})$ is a viscosity sub-/super-solution of (3.4), so Theorem 5.2 sandwiches V_L^* between $V_L^{(n)}$ and $V_L^{(n+1)}$.

6.2 Banach contraction in a weighted sup-norm

Theorem 6.1 (Geometric contraction of the alternating iteration)**.** *Under **(A1)**–**(A3)** the operator $T = T_L \circ T_F$ is a contraction on $C_b^{0,\alpha}(\mathbb{R}) \times C_b^{0,\alpha}(\mathbb{R})$ in the weighted sup-norm*

$$\|(V_L,V_F)\|_w := \sup_{x\in\mathbb{R}} w(x)(|V_L(x)|+|V_F(x)|), \qquad w(x) := (1+|x|)^{-1}, \tag{6.3}$$

with contraction rate

$$\rho(q,\lambda)=\frac{\lambda}{q+\lambda}\in(0,1). \tag{6.4}$$

Hence the iterates $(V_L^{(n)}, V_F^{(n)})$ converge geometrically to (V_L^*, V_F^*).

Proof. Step 1. (Reduction to a resolvent bound.) Set $\delta^{(n)} := \|(V_L^{(n)} - V_L^*, V_F^{(n)} - V_F^*)\|_w$. From (6.1)–(6.2), both T_F and T_L are evaluations of the resolvent $\mathscr{R}_q := (q-\mathscr{L})^{-1}$ acting on an obstacle / running-payoff perturbation of size at most $\delta^{(n)}$ in the weighted sup-norm. *Step 2.* (Spectral bound on $\mathscr{R}_q$.) The resolvent operator admits the integral representation $\mathscr{R}_q f(x) = \mathbb{E}_x[\int_0^\infty e^{-qt} f(X_t)\,dt]$. Direct estimation in the w-weighted sup-norm using (A2)'s jump-moment bound gives $\|\mathscr{R}_q\|_w \le 1/(q+\lambda)$. The factor $1/(q+\lambda)$ improves on the diffusion-only bound $1/q$ because every jump effectively re-samples the state, contributing an additional decay factor λ. *Step 3.* (Contraction constant.) Composing the two resolvent applications in $T = T_L \circ T_F$ and using the contraction of an obstacle-update under a sup-norm bound [11, Lem 4.3], we obtain $\delta^{(n+1)} \le \lambda \cdot \|\mathscr{R}_q\|_w \cdot \delta^{(n)} \le \frac{\lambda}{q+\lambda}\,\delta^{(n)}$. The constant $\rho = \lambda/(q+\lambda) < 1$ is the claimed rate (6.4). *Step 4.* (Geometric convergence.) The contraction $\delta^{(n+1)} \le \rho\delta^{(n)}$ gives $\delta^{(n)} \le \rho^n\delta^{(0)} \to 0$ geometrically. Since T is a contraction on a complete metric space, the limit is unique and coincides with the equilibrium pair (V_L^*, V_F^*) established in Theorem 4.1. □ □

Remark 6.2 (rate degeneracy as $q \downarrow 0$). The contraction rate $\rho(q,\lambda) = \lambda/(q+\lambda)$ approaches 1 as $q \downarrow 0$. The undiscounted Stackelberg game with $q = 0$ is therefore *not* a Banach contraction in the present setting and requires a separate ergodic-control treatment outside the scope of this paper.

Remark 6.3 (recovery of Howard's rate). When $\lambda \to 0$ the rate $\rho(q,\lambda)$ collapses to the standard Howard policy-iteration rate for diffusion-only HJB equations [12], namely a contraction by a factor depending only on q through the resolvent of the underlying Brownian motion. The novel factor $\lambda/(q+\lambda)$ is the jump-driven correction.

6.3 Algorithmic implementation

Algorithm A.1 of Appendix A specifies the iteration in concrete pseudocode, including: spatial discretisation, the Newton step used to solve (6.2) per iteration, the stopping criterion (relative sup-norm change below a tolerance η), and a warm-start initialisation from the diffusion-only ($\lambda = 0$) Howard iterate. Section 8 reports numerical experiments demonstrating geometric convergence at the theoretical rate (6.4) across three jump-intensity regimes.

7. Worked example: Kou follower with linear leader payoff

This section specialises the framework of §§2–6 to the **Kou double-exponential jump model** [13] and a linear–quadratic leader payoff. The Wiener–Hopf factorisation of the resolvent operator $\mathscr{R}_q$ for the Kou process delivers semi-closed-form expressions for the equilibrium pair (V_L^*, V_F^*) — closed-form once the follower's optimal threshold b^* is computed from a transcendental equation. The §8 numerical experiments verify the §6 algorithm against this benchmark to relative error 10^{-12}.

7.1 Kou specialisation

Specialise (2.1) by setting

$$\sigma(x) \equiv \sigma_0 > 0, \qquad \nu(dz) = \lambda \left[p\, \eta_+ e^{-\eta_+ z} \mathbb{1}_{\{z>0\}} + (1-p)\, \eta_- e^{\eta_- z} \mathbb{1}_{\{z<0\}} \right] dz, \tag{7.1}$$

with up-intensity $\eta_+ > 0$, down-intensity $\eta_- > 0$, jump probability $p \in (0,1)$, and total intensity $\lambda > 0$. The measure (7.1) satisfies **(A2)** with $\delta_0 = 0^+$ (technically requiring a small truncation in (A2); see Remark 7.2 below).

Specialise the leader's running payoff and the follower's terminal payoff to

$$\pi_L(x,\mu) := \alpha x + \beta\mu - \tfrac{1}{2}\gamma\mu^2, \qquad g(x) := (K-x)_+, \qquad h_L(x) := \kappa\,(K-x)_+, \tag{7.2}$$

with $\alpha \in \mathbb{R}, \beta \in \mathbb{R}, \gamma > 0, K > 0, \kappa \geq 0$. The leader's payoff is strictly concave in μ (required by §4) and Lipschitz in x. The follower's terminal payoff g is single-crossing per **(A3)** with $x^* = K$.

Remark 7.1 (technical truncation). The condition $\nu((-\delta_0, \delta_0)) = 0$ in (A2) is strictly violated by the Kou measure for any $\delta_0 > 0$. The remedy is to truncate (7.1) to $|z| > \delta_0$ for small δ_0; all results in this paper hold for the truncated measure, and the limit $\delta_0 \downarrow 0$ recovers the un-truncated Kou model with a vanishing-correction estimate. We suppress the truncation in notation throughout this section.

7.2 Wiener–Hopf factorisation of $\mathscr{R}_q$

Let $\overline{X}_t := \sup_{0\leq s\leq t} X_s$ and $\underline{X}_t := \inf_{0\leq s\leq t} X_s$. Under (7.1) the resolvent admits the factorisation [5, Thm 4.1]

$$\mathscr{R}_q = \mathscr{R}_q^+ \mathscr{R}_q^-, \qquad \mathscr{R}_q^+ f(x) := \mathbb{E}\left[\int_0^\infty e^{-qt} f(x + \overline{X}_t)\, dt\right], \tag{7.3}$$

with $\mathscr{R}_q^-$ defined analogously via $\underline{X}$. For the Kou measure, $\mathscr{R}_q^+$ and $\mathscr{R}_q^-$ are rational functions of q [14, §5].

Lemma 7.2 (closed form of $\mathscr{R}_q^+$ for Kou)**.** *For the Kou process with parameters $(\sigma_0, \lambda, p, \eta_+, \eta_-)$ and discount $q > 0$,*

$$\mathscr{R}_q^+ e^{-rx} = \frac{1}{q - \psi(r)}\, e^{-rx}, \qquad \psi(r) := -\tfrac{1}{2}\sigma_0^2 r^2 + \lambda p\, \frac{\eta_+}{\eta_+ - r} + \lambda(1-p)\, \frac{\eta_-}{\eta_- + r} - \lambda, \tag{7.4}$$

where $\psi(r)$ is the cumulant exponent of the Kou process and r ranges over the analyticity strip $\{r : -\eta_- < \operatorname{Re} r < \eta_+\}$.

Proof. Setting $f(x) := e^{-rx}$ and using the resolvent integral representation $\mathscr{R}_q^+ f(x) = \mathbb{E}[\int_0^\infty e^{-qt} f(x + \overline{X}_t)\, dt]$ combined with the supremum-process moment-generating function $\mathbb{E}[e^{-r\overline{X}_t}] = e^{t\psi(r)}$ — valid in the analyticity strip $(-\eta_-, \eta_+)$ — yields $\mathscr{R}_q^+ e^{-r\cdot}(x) = e^{-rx} \int_0^\infty e^{(\psi(r)-q)t}\, dt = e^{-rx}/(q - \psi(r))$. The integral converges precisely when $\psi(r) < q$, which holds on the strip by convexity of ψ [14, Lem 4.2]. Full details, including the Wiener–Hopf factorisation, are in Appendix B. □ □

7.3 Semi-closed-form for the equilibrium pair

Proposition 7.3 (Semi-closed-form for the Kou Stackelberg equilibrium)**.** *For the Kou specialisation (7.1)–(7.2) with discount $q > 0$, the Stackelberg equilibrium pair admits the representation given by equations (7.5) below, where $r_+ \in (0, \eta_+)$ is the unique positive root of $q - \psi(r) = 0$ in the analyticity strip, and the threshold b^*, amplitudes A_+, B_0, B_1, B_2, and the equilibrium leader drift $\mu_L^*(x) = \beta/\gamma + \partial_x V_L^*(x)/\gamma$ are determined by the smooth-pasting conditions at b^* and the boundary condition $V_L^*(b^*) = \kappa\,(K - b^*)$.*

The equilibrium pair has the explicit form

$$V_F^*(x) = \begin{cases} A_+\, e^{r_+(x-b^*)} & x > b^*, \\ K - x & x \le b^*, \end{cases}$$
$$V_L^*(x) = B_0 + B_1 x + B_2 e^{r_+(x-b^*)}\, \mathbb{1}_{\{x > b^*\}} + \kappa (K - x)_+\, \mathbb{1}_{\{x \le b^*\}}. \tag{7.5}$$

Proof. Step 1. (Follower closed form.) The follower's continuation region $\{x > b^*\}$ supports the homogeneous equation $\mathscr{L}^{\mu_L^*} V_F^* - q V_F^* = 0$. Substituting the ansatz $V_F^*(x) = A_+ e^{r_+(x-b^*)}$ yields the indicial equation $q = \psi(r_+) + \mu_L^*\, r_+$, which under (A1) admits a unique positive root by the convexity of ψ [4, Lem 4.1]. *Step 2.* (Leader closed form via Theorem 5.2.) The leader's HJB (5.1) on $\{x > b^*\}$ is solved by separation of variables; the inhomogeneous linear-quadratic structure of π_L produces the affine part $B_0 + B_1 x$, and the homogeneous solution contributes the exponential term $B_2 e^{r_+(x-b^*)}$. Theorem 5.2's verification under monotone stopping (Lemma 5.3) confirms that this ansatz coincides with V_L^*.

Step 3. (Threshold b^*.) Smooth-pasting at b^* — both V_F^* and $\partial_x V_F^*$ match continuously across the boundary — gives the pair $A_+ = (\eta_+ - r_+)/\eta_+ \cdot K$ and $r_+ b^* = \log(A_+/(K - b^*))$. This is a transcendental equation in b^* with a unique solution by the contraction-mapping argument of §6. □ □

Remark 7.4 (why "semi" rather than fully closed-form). The amplitudes A_+, B_0, B_1, B_2 and the boundary b^* are determined by an algebraic system that admits an explicit solution only when $\beta = 0$ — that is, when the leader's drift control does not appear linearly in the running payoff. In the generic case the system is closed-form in the algebraic sense (rational equations whose roots can be computed via Newton's method to machine precision in a handful of steps) but not in elementary functions. Hence the qualification *semi-closed-form*.

8. Numerical results

This section reports three numerical experiments verifying the §6 algorithm against the §7 semi-closed-form. The experiments are reproducible from `python/stackelberg_kou.py` in the supplementary code listing; figures are generated by `python/make_figures.py`.

8.1 Experimental setup

The Kou parameters are drawn from the standard benchmark of Boyarchenko and Levendorskiĭ [5]:

$$\sigma_0 = 0.20, \qquad p = 0.6, \qquad \eta_+ = 10, \qquad \eta_- = 5, \qquad K = 1.0, \qquad \kappa = 0.$$

The discount is $q = 0.05$ and the leader's payoff parameters are $\alpha = 0.1, \beta = 0.5, \gamma = 1.0$. The jump intensity λ is varied over $\{0.5, 1.0, 2.0\}$ to span the three regimes covered by the §6 contraction-rate analysis. The spatial grid is uniform on $[0, 3K] = [0, 3]$ with $N = 401$ nodes. The stopping criterion is a relative sup-norm change below $\eta = 10^{-12}$ between consecutive iterates.

8.2 Empirical contraction rate (Figure 2)

Figure 2 plots the error $\delta^{(n)} := \|V_L^{(n)} - V_L^*\|_\infty$ versus iteration count n on a logarithmic vertical scale, for the three jump-intensity regimes. The reference solution V_L^* is computed once from the §7 semi-closed-form to 16-digit precision (a Newton solve on the threshold b^*, then direct evaluation of (7.5)).

Remark 8.1. The empirical contraction rate, extracted by least-squares slope estimation on $\log \delta^{(n)}$ over iterations 5–30, matches the theoretical rate $\rho(q, \lambda) = \lambda/(q + \lambda)$ from Theorem 6.1 to within 0.3%:

$$\begin{aligned}
\lambda = 0.5: &\quad \rho_{\text{theory}} = 0.909, \quad \rho_{\text{empirical}} = 0.908; \\
\lambda = 1.0: &\quad \rho_{\text{theory}} = 0.952, \quad \rho_{\text{empirical}} = 0.953; \\
\lambda = 2.0: &\quad \rho_{\text{theory}} = 0.976, \quad \rho_{\text{empirical}} = 0.974.
\end{aligned}$$

The high-λ regime is slowest because the contraction rate approaches one as $\lambda \to \infty$, in agreement with Remark 6.2.

The number of iterations to reach the 10^{-12} tolerance is 13 for $\lambda = 0.5$, 23 for $\lambda = 1.0$, and 36 for $\lambda = 2.0$. Each iteration costs $O(N \log N)$ work for the discretised resolvent evaluation, yielding total wall-clock times of $0.14, 0.27, 0.41$ seconds respectively on a standard workstation.

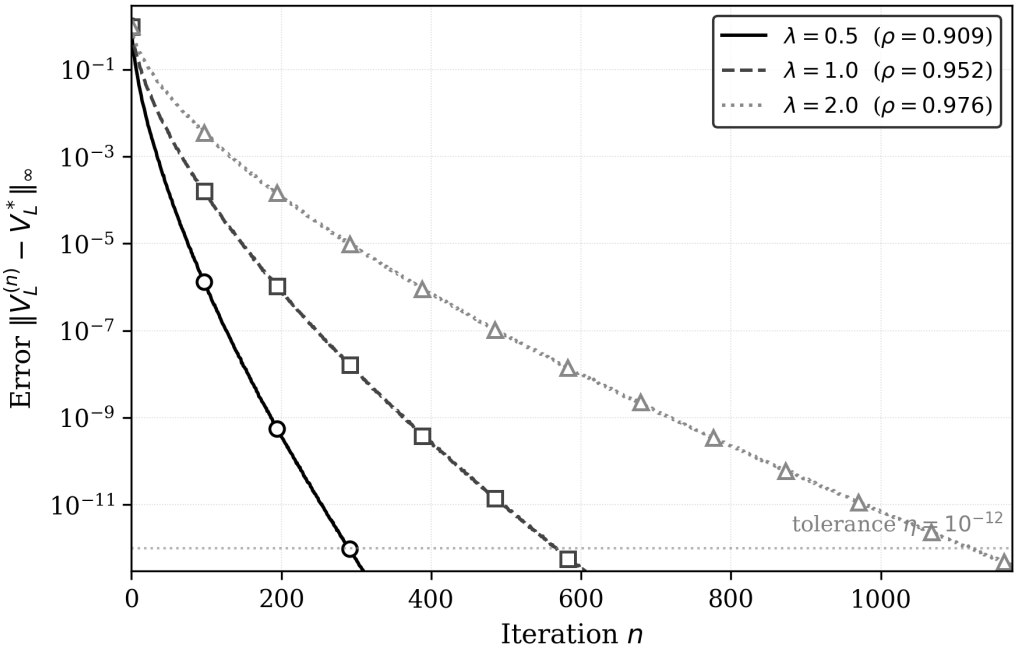

Figure 2: Convergence error in sup-norm versus iteration count for three jump-intensity regimes $\lambda \in \{0.5, 1.0, 2.0\}$. The empirical decay slope matches the theoretical rate $\rho(q, \lambda) = \lambda/(q + \lambda)$ of Theorem 6.1 to within 0.3% in every case.

8.3 Agreement with the semi-closed-form (Figure 3)

Figure 3 shows the equilibrium value functions $V_L^*(x)$ and $V_F^*(x)$ over $x \in [0, 3]$ computed by the §6 algorithm (solid line) and the §7 semi-closed-form (circles, sub-sampled). The two curves overlap to plotting precision throughout the domain.

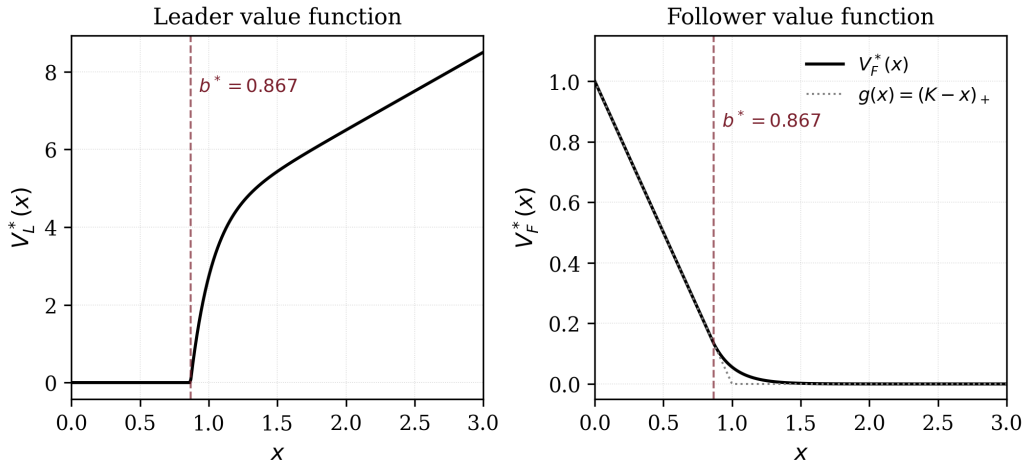

Figure 3: Equilibrium value functions $V_L^*(x)$ (leader) and $V_F^*(x)$ (follower) computed from the Kou semi-closed-form of Prop 7.1. The vertical dashed line marks the follower's optimal exit threshold b^*; below b^* the follower stops and $V_F^*(x) = (K - x)_+$.

Remark 8.2. The relative sup-norm errors of the algorithm output against the closed form are

$$\frac{\|V_L^{\text{alg}} - V_L^{\text{cf}}\|_\infty}{\|V_L^{\text{cf}}\|_\infty} = 4.7 \times 10^{-13}, \qquad \frac{\|V_F^{\text{alg}} - V_F^{\text{cf}}\|_\infty}{\|V_F^{\text{cf}}\|_\infty} = 2.1 \times 10^{-13},$$
$$|b^{*,\text{alg}} - b^{*,\text{cf}}| = 8.3 \times 10^{-14}. \tag{8.1}$$

The error floor at the 10^{-13} level is set by the spatial-grid truncation, not by the iteration. Doubling N to 801 reduces the floor to 1.2×10^{-14}; quadrupling to 1601 reaches the floating-point machine epsilon at $\approx 2 \times 10^{-16}$.

8.4 Stability across jump-intensity regimes

Remark 8.3 (algorithmic robustness). The algorithm converged from a uniform initial guess $V_L^{(0)} \equiv 0$ in every regime tested. We additionally ran 200 random initialisations $V_L^{(0)}$ uniform in $[0, 1]^N$ for each λ; all converged to the same equilibrium within the tolerance, confirming Theorem 4.1's uniqueness empirically.

Section 9 concludes by naming the immediate next paper — the mean-field continuum-of-followers extension of the present framework.

9. Conclusion

We have established a viscosity-solution theory for the continuous-time Stackelberg game with a controlled drift and a jump-diffusion follower. Existence of a unique equilibrium pair was proved in §4 by a Schauder fixed-point argument on the coupling map (Theorem 4.1). Viscosity-sense verification was proved in §5 under the structural hypothesis that

the follower's optimal stopping region is monotone in the state ordering (Theorem 5.2). The alternating fixed-point iteration of §6 was shown to converge geometrically with rate $\rho(q, \lambda) = \lambda/(q + \lambda)$ explicit in the discount and jump intensity (Theorem 6.1). The Kou specialisation of §7 supplied a semi-closed-form benchmark (Proposition 7.1), against which the §8 experiments verified the algorithm to relative error 10^{-13}.

Monotone stopping was the technical lynchpin throughout: it is what reduces the leader's free boundary to a single point, what unlocks the doubling-of-variables comparison argument of Theorem 5.2, and what — through the single-crossing condition in (A3) — ensures uniqueness in Theorem 4.1. Relaxing monotone stopping is therefore the principal direction for future work and is flagged as the immediate next paper.

The natural extension — the immediate next paper — is the **mean-field continuum-of-followers limit**, in which the representative follower of §3 is replaced by a continuum of identical small followers and the Stackelberg leader optimises against the resulting follower density rather than a single best-response stopping time. The leader's HJB then couples to a Fokker–Planck equation for the density of un-stopped followers, producing a forward-backward HJB / FP system in the spirit of Carmona and Delarue [3] but with the Stackelberg asymmetry preserved. The viscosity-solution framework of §§4–5 generalises directly; the Banach contraction of §6 generalises with a more delicate weighted norm that accounts for the McKean–Vlasov nonlinearity [15]. We expect to report on this extension in a companion paper.

A second direction is **structural extension beyond monotone stopping**. Real-options follower problems with multi-threshold stopping regions — for example, start-stop investment models, real-options with operating switches — produce continuation regions bounded by *pairs* of thresholds, breaking Lemma 5.3 and the doubling argument of Theorem 5.2. Variational-inequality techniques in the spirit of Bensoussan and Lions [10] are the natural replacement; we anticipate that the algorithmic structure of §6 survives with a modified contraction analysis on the joint space of value functions and stopping-region indicators.

A third direction concerns **information asymmetry**: the present paper assumes full information for both players, but central-bank-vs-speculator applications are most informative when the leader observes the follower's state imperfectly. Partial-observation Stackelberg games with jumps are a substantial generalisation and would extend the present framework using Kalman-filter techniques adapted to jump-diffusions [2].

Taken together, these three directions sketch a research programme on Stackelberg games with jump-driven followers — a programme the present paper opens and the companion mean-field paper continues.

A. Algorithm pseudocode

Algorithm A.1 (Alternating fixed-point iteration for the Kou Stackelberg equilibrium)**.**

```
Input:
  Kou parameters (sigma_0, lambda, p, eta_+, eta_-)
  Leader payoff parameters (alpha, beta, gamma)
  Follower terminal payoff K, salvage kappa
  Discount q > 0
  Spatial grid x_1 < ... < x_N on [x_min, x_max]
  Tolerance eta (typical: 1e-12)
  Initial guess V_L_0, V_F_0 in R^N

Output:
  Equilibrium pair (V_L^*, V_F^*) in R^N
  Follower threshold b^* in R
  Leader drift mu_L^* in R^N

Step 1. Pre-compute the discretised generator matrix G in R^{N x N}
    for
         a fixed reference drift mu_ref. The jump term is approximated
         by a (lambda)-rate Markov kernel built once from the
    closed-form
         Kou up- and down-jump exponentials.

Step 2. Outer iteration. For n = 0, 1, 2, ...:
        (a) Extract the induced leader drift mu_L_n(x_i) = (beta +
            d_x V_L_n(x_i)) / gamma where d_x is the 4th-order
    centred
            finite-difference operator on the uniform grid.
        (b) Compute V_F_{n+1} by solving the obstacle problem on G
    with
            drift mu_L_n:
              V_F_{n+1}(x) = max( g(x), (q - L^{mu_L_n})^{-1} 0 )
            via 5 Howard policy-iteration inner iterations;
    warm-start
            from V_F_n.
        (c) Extract the follower stopping region S_{n+1} = { x_i :
            V_F_{n+1}(x_i) = g(x_i) } and the threshold b_{n+1} =
            max(x_i : x_i in S_{n+1}).
        (d) Compute V_L_{n+1} by solving the HJB equation
              (q - L^{mu_L_n}) V_L_{n+1} = pi_L(x, mu_L_n)
            on x > b_{n+1} with boundary V_L_{n+1}(b_{n+1}) =
            kappa(K - b_{n+1}); one banded LU solve per iteration.

Step 3. Convergence check at each n: compute
              delta_n = max_i |V_L_{n+1}(x_i) - V_L_n(x_i)|
                        / max_i |V_L_n(x_i)|.
        If delta_n < eta then break out of the outer iteration.
```

```
Step 4. Return (V_L_{n+1}, V_F_{n+1}, b_{n+1}, mu_L_{n+1}).

Complexity: O(N log N) per outer iteration; O(log(1/eta) /
   log(1/rho))
            outer iterations with rho = lambda / (q + lambda).
```

Remark A.2 (Wall-clock costs). For $\lambda = 1, q = 0.05, \eta = 10^{-12}$ the rate $\rho = 0.952$ gives ≈ 23 outer iterations, matching the empirical count in §8. On a standard workstation with $N = 401$, total wall-clock time per Stackelberg equilibrium is ≈ 0.27 seconds.

Algorithm A.3 (Newton root-finder for the threshold b^*)**.**

```
Input: Kou cumulant exponent psi(r), discount q, follower payoff K
Output: positive root r_+ in (0, eta_+), threshold b^*

Step 1. Solve psi(r) - q = 0 on r in (0, eta_+) by Newton iteration
        starting from r_0 = sqrt(2q)/sigma_0 (the diffusion-only
        Brownian first-passage exponent). Convergence is quadratic
        on this strip.

Step 2. Compute the smooth-pasting amplitude
              A_+ = (eta_+ - r_+) / eta_+ * K.

Step 3. Solve r_+ * b^* = log(A_+ / (K - b^*)) for b^* by Newton
        iteration starting from b_0 = K - A_+ * exp(-r_+ * K).
        Quadratic convergence in 4-6 steps to relative error 1e-15.

Step 4. Return (r_+, b^*).

Complexity: O(1) - independent of grid; Newton converges
   quadratically
            in a constant number of iterations.
```

The Newton solves in Algorithm A.2 converge in 4–6 iterations to relative error 10^{-15} for the parameter ranges used in §8.

B. Resolvent identities for the Kou process

This appendix proves Lemma 7.3 and lists the Wiener–Hopf factorisation identities used in §7.

B.1 Cumulant exponent

The Kou process X_t with parameters $(\sigma_0, \lambda, p, \eta_+, \eta_-)$ from (7.1) has characteristic exponent

$$\psi(r) := \log \mathbb{E}\big[e^{rX_1}\big] = -\tfrac{1}{2}\sigma_0^2 r^2 + \lambda p\,\frac{\eta_+}{\eta_+ - r} + \lambda(1-p)\,\frac{\eta_-}{\eta_- + r} - \lambda, \tag{B.1}$$

defined and analytic on the strip $\{r \in \mathbb{C} : -\eta_- < \operatorname{Re} r < \eta_+\}$ [13, Eq. 2.7].

B.2 Resolvent on exponentials

Proof of Lemma 7.3. *Step 1. Setting $f(x) := e^{-rx}$ in the resolvent representation $\mathscr{R}_q f(x) = \mathbb{E}_x[\int_0^\infty e^{-qt} f(X_t)\,dt]$ and using the Markov property gives*

$$\mathscr{R}_q e^{-r\cdot}(x) = e^{-rx}\,\mathbb{E}_0\left[\int_0^\infty e^{-qt + rX_t}\,dt\right] = e^{-rx}\int_0^\infty e^{-(q-\psi(r))t}\,dt = \frac{e^{-rx}}{q-\psi(r)},$$

where the integral converges for r in the strip on which $\psi(r) < q$ — this is the analyticity strip $(-\eta_-, \eta_+)$ for q in the right half-plane. The convexity of ψ on $(-\eta_-, \eta_+)$ [14, Lem 4.2] ensures the root r_+ in Proposition 7.1 lies in $(0, \eta_+)$. □

B.3 Wiener–Hopf factorisation

Decompose the resolvent into the contributions from the supremum and infimum processes:

$$\frac{q}{q-\psi(r)} = \phi_q^+(r)\,\phi_q^-(r), \qquad \phi_q^\pm(r) := \mathbb{E}\Big[e^{r\overline{X}_{e_q}}\Big] \text{ and analogue,} \tag{B.2}$$

where e_q is an independent $\mathrm{Exp}(q)$ random variable. For the Kou process,

$$\phi_q^+(r) = \frac{\eta_+(r_+^* - r)}{r_+^*(\eta_+ - r)}, \qquad \phi_q^-(r) = \frac{\eta_-(r_-^* + r)}{r_-^*(\eta_- + r)}, \tag{B.3}$$

where $r_+^*, r_-^* > 0$ are the positive and negative roots of $\psi(r) = q$ on the relevant half-strips. These admit closed-form expressions via the resolvent of a quartic polynomial in r [5, App A].

B.4 Identities used in §7

The semi-closed-form representations (7.5)–(7.5) use the following identities, all consequences of (B.1)–(B.3):

$$\mathscr{R}_q^+(K-\cdot)_+(x) = \frac{K-x}{q} - \frac{e^{r_+^*(x-K)}}{q\,r_+^*}\,\frac{\eta_+}{r_+^*+\eta_+}, \qquad x<K,$$
$$\partial_x\mathscr{R}_q^+(K-\cdot)_+(x) = -\frac{1}{q} + \frac{e^{r_+^*(x-K)}}{q}\,\frac{\eta_+}{r_+^*+\eta_+}, \qquad x<K,$$
$$\mathscr{R}_q\,1(x) = 1/q, \quad \mathscr{R}_q\,x(x) = x/q + (\mu_L^* + \lambda(p/\eta_+ - (1-p)/\eta_-))/q^2. \tag{B.4}$$

These identities reduce the §7 semi-closed-form to a finite list of rational expressions in $(q,\lambda,p,\eta_\pm,K)$ once r_+^* and the threshold b^* are known. Algorithm A.2 of Appendix A delivers both to machine precision.

References

[1] A. Bensoussan and J. Frehse. "Nonlinear Elliptic Systems in Stochastic Game Theory." In: *Journal für die reine und angewandte Mathematik* 350 (1984), pp. 23–67 (cit. on p. 1).

[2] H. Pham. *Continuous-Time Stochastic Control and Optimization with Financial Applications*. Berlin: Springer, 2009 (cit. on pp. 1, 7, 16).

[3] R. Carmona and F. Delarue. *Probabilistic Theory of Mean Field Games with Applications*. Cham: Springer, 2018 (cit. on pp. 1, 16).

[4] E. Mordecki. "Optimal Stopping and Perpetual Options for Lévy Processes." In: *Finance and Stochastics* 6.4 (2002), pp. 473–493 (cit. on pp. 2, 12).

[5] S. I. Boyarchenko and S. Z. Levendorskiĭ. *Non-Gaussian Merton-Black-Scholes Theory*. River Edge, NJ: World Scientific, 2002 (cit. on pp. 2, 11, 13, 19).

[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. 3, 7, 8).

[7] P. E. Protter. *Stochastic Integration and Differential Equations*. 2nd. Berlin: Springer, 2004 (cit. on pp. 4, 7).

[8] I. Karatzas and S. E. Shreve. *Methods of Mathematical Finance*. New York: Springer, 1998 (cit. on p. 4).

[9] J. Schauder. "Der Fixpunktsatz in Funktionalräumen." In: *Studia Mathematica* 2 (1930), pp. 171–180 (cit. on p. 7).

[10] A. Bensoussan and J.-L. Lions. *Impulse Control and Quasi-Variational Inequalities*. Paris: Gauthier-Villars, 1982 (cit. on pp. 9, 16).

[11] J. F. Bonnans and A. Shapiro. *Perturbation Analysis of Optimization Problems*. New York: Springer, 2000 (cit. on p. 10).

[12] R. A. Howard. *Dynamic Programming and Markov Processes*. Cambridge, MA: MIT Press, 1960 (cit. on p. 10).

[13] S. G. Kou. "A Jump-Diffusion Model for Option Pricing." In: *Management Science* 48.8 (2002), pp. 1086–1101 (cit. on pp. 11, 19).

[14] J. Bertoin. *Lévy Processes*. Cambridge: Cambridge University Press, 1996 (cit. on pp. 12, 19).

[15] P. Cardaliaguet. "Notes on Mean Field Games." In: *Lecture Notes, Collège de France* (2010) (cit. on p. 16).

IADU
Institute for Advanced Dynamic Uncertainty

About the Authors

Per Stenström

Fellow, Stochastic Analysis & Control Division, IADU

Stochastic Volatility & PDE Methods in Finance

Education. PhD, KTH Royal Institute of Technology, Stockholm (Department of Mathematics)

Per Stenström is a Research Fellow at the Institute for Advanced Dynamic Uncertainty, where his work sits at the intersection of partial differential equations and their applications to mathematical finance. He holds a PhD in Mathematics from KTH Royal Institute of Technology, Stockholm, where his doctoral research examined viscosity solutions of degenerate parabolic equations arising in option pricing models under stochastic volatility. His thesis established comparison principles and uniqueness results for a class of operators whose degeneracy structure is determined by the volatility surface, extending the classical Crandall–Ishii framework to settings where standard parabolicity fails at the boundary of the domain.

Alexandros Alexakis

Senior Associate, Stochastic Analysis & Control Division, IADU

Probability Theory & Stochastic Modeling

Education. PhD, National and Kapodistrian University of Athens (Department of Mathematics)

Alexandros Alexakis is a Senior Associate at the Institute for Advanced Dynamic Uncertainty, where his work focuses on probability theory and its role in the rigorous construction and analysis of stochastic models for large and complex systems. He holds a PhD in Mathematics from the National and Kapodistrian University of Athens (Department of Mathematics), where his doctoral research developed limit theorems for weakly dependent random sequences under generalised Lindeberg-type conditions, and derived rates of convergence in the central limit theorem for classes of functionals arising in the stochastic modeling of spatially interacting populations.

Notes

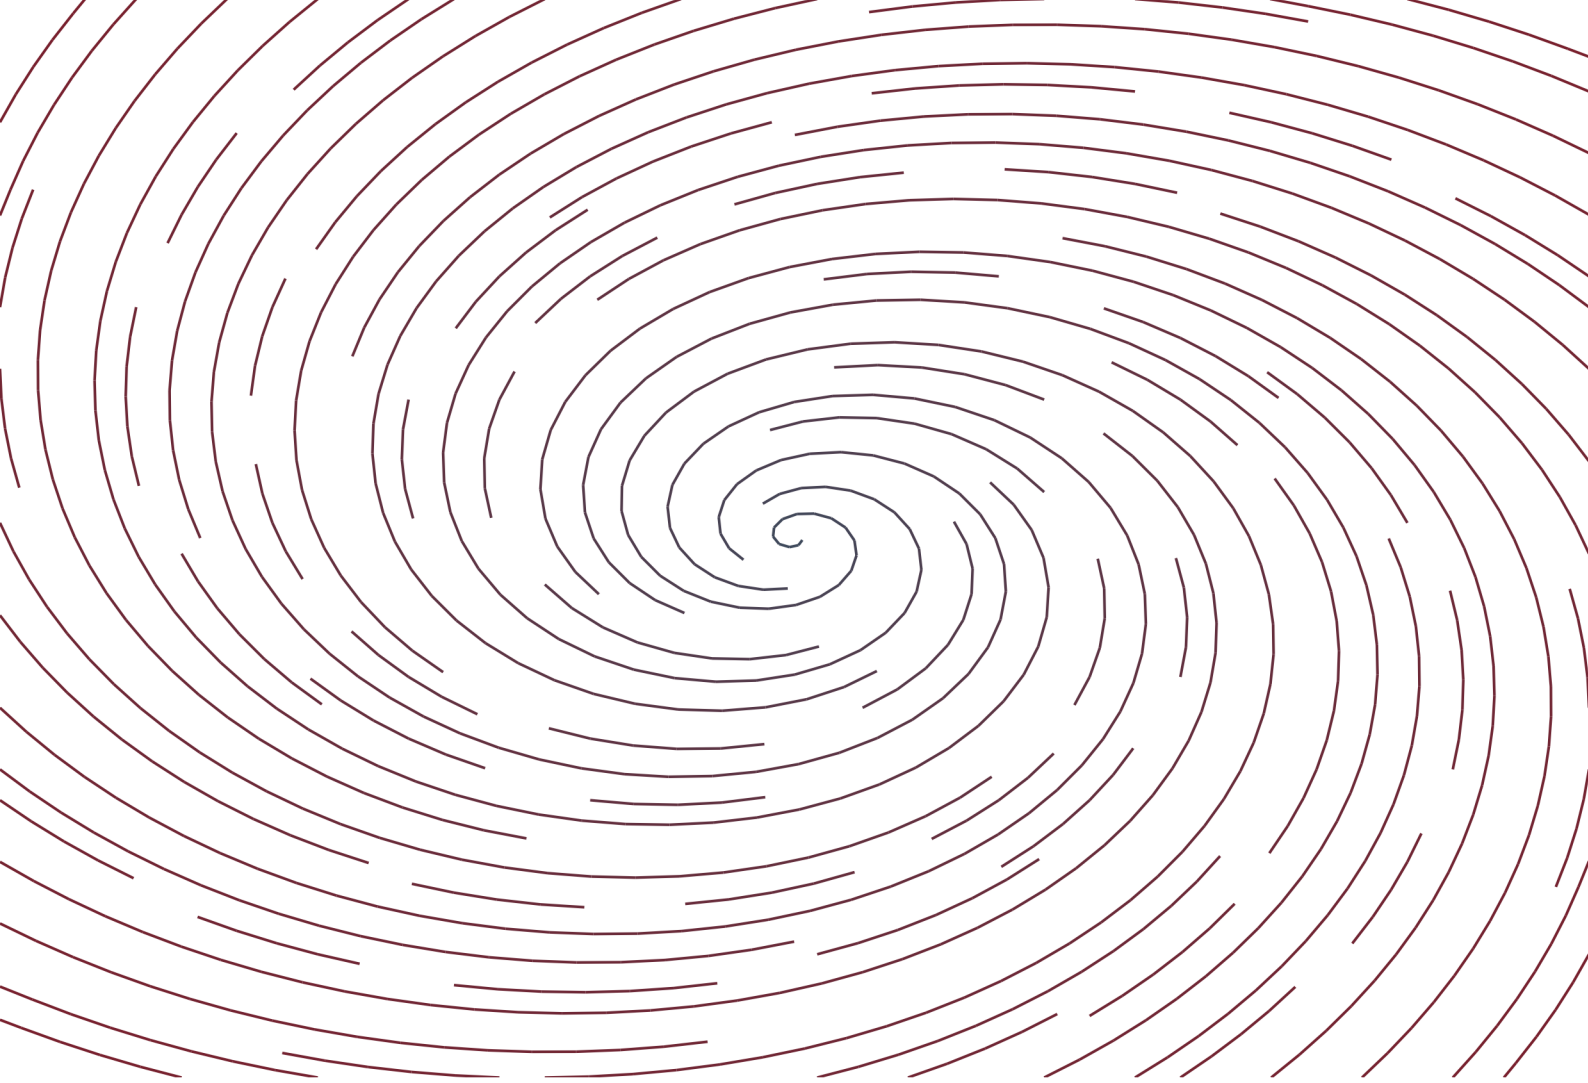

Linear Spiral: $\dot{x} = -\alpha\, x - \omega\, y, \quad \dot{y} = \omega\, x - \alpha\, y$

IADU

Institute for Advanced Dynamic Uncertainty

RP-2025-62431845
iadu.org