

PREPRINT

Stochastic Stackelberg Games with a Jump-Diffusion Follower

Existence, Verification, and a Fixed-Point Algorithm

Supervised by Dr. T. Zamrik

13-APR-2025



Sample paths of the Ornstein–Uhlenbeck process: $dX_t = \kappa(\theta - X_t)dt + \sigma dW_t$, $\kappa = 3.68$, $\theta = 0$, $\sigma = 0.72$, $X_0 = -1.25$



PR-2026-46511496
zamrik.com



Zamrik Quant Research

The research arm of Zamrik Quant Lab

Copyright

© 2025 T. Zamrik. All rights reserved. Published by Zamrik Quant Research. No part of this document may be reproduced, stored, distributed or adapted, in whole or in part, by any means, without the prior written permission of the publisher, except for brief quotation with full citation. Requests for permission: info@zamrik.com.

Suggested Citation

Optimal Control & Decision Desk, Zamrik Quant Research (2025). ‘Stochastic Stackelberg Games with a Jump-Diffusion Follower: Existence, Verification, and a Fixed-Point Algorithm’. *Zamrik Quant Research Preprint* **PR-2026-46511496**. Available at <https://zamrik.com/research-items/pr-2026-46511496/>.

About Zamrik Quant Research

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

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

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

Legal Notice

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



PREPRINT

Stochastic Stackelberg Games with a Jump-Diffusion Follower

Existence, Verification, and a Fixed-Point Algorithm

Supervised by Dr. T. Zamrik¹

13-APR-2025

PR-2026-46511496

Abstract. We study a continuous-time Stackelberg game in which a leader sets the drift of a one-dimensional jump-diffusion through a feedback rule and a follower, who observes the rule, chooses when to exit. In the feedback sense, each rule a best response to the other, the equilibria are the solutions of a coupled system: the leader's Hamilton–Jacobi–Bellman (HJB) equation off the follower's exit region, and the follower's obstacle problem under the leader's drift. We prove existence of an equilibrium by a Schauder–Tychonoff argument, a half-line structure of the exit region for a put-type exit payoff under a non-increasing drift, and a comparison theorem for the leader's equation when the exit region is a half-line. For the alternating fixed-point algorithm we show that each half-step, computed by splitting off the jump term, is a sup-norm contraction with rate at most $\lambda/(q+\lambda)$, and that in the half-line case the outer iteration is a scalar map on the follower's threshold whose rate is not fixed by q and λ . For double-exponential (Kou) jumps the follower's best response to a constant drift is in closed form through the Wiener–Hopf factors of the resolvent paper of the Stochastic Analysis Desk: the threshold is $K - 1/\beta_1 - 1/\beta_2 + 1/\eta_-$. A grid implementation of the two half-steps reproduces this closed form at second order, and the follower sweeps contract at measured rates 0.899, 0.947 and 0.973 against the bounds 0.909, 0.952 and 0.976.

PR-2026-46511496 · Optimal Control & Decision Desk · Preprint · Published
13-APR-2025 · Reading time: about 28 min

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 follower best-responds to the leader's announced drift rule, and the leader's rule must be optimal given how the follower exits. This is a continuous-time Stackelberg game; the follower's right to stop introduces a free boundary into the leader's problem, and the jumps make every equation non-local, so the viscosity-solution techniques developed for diffusion-only games [1] need adapting.

The literature has approached this problem from three sides. Two-player zero-sum stochastic differential games, beginning with Bensoussan and Frehse [1], and the control of jump-

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



diffusions, treated by Pham [2], cover simultaneous optimisation; the Stackelberg asymmetry, in which one player's rule is announced and the other responds to it, falls outside that framework. Mean-field Stackelberg games [3] address a continuum of followers but presuppose a regularity infrastructure (the master equation, McKean–Vlasov dynamics) that obscures the single-follower analysis. Perpetual-option pricing under Lévy noise, from Mordecki [4] and Boyarchenko and Levendorskiĭ [5], supplies closed forms for the follower's stopping rule but treats the drift as exogenous. For those closed forms we use the resolvent paper of the Stochastic Analysis Desk [6], whose Wiener–Hopf factors and perpetual-put theorem are quoted here as stated there.

Notation 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 through a feedback rule $\mu_L(x)$. The follower chooses a stopping time $\tau \in \mathcal{T}$ adapted to $(\mathcal{F}_t)_{t \geq 0}$; $\mathbf{e}(q)$ is an exponential time of rate q independent of X .

This paper establishes four results.

Claim 1 (Theorem 1). Under the standing assumptions (A1)–(A3) and a continuity hypothesis (H) on the best responses, a feedback Stackelberg equilibrium exists, and its value functions (V_L^*, V_F^*) solve the coupled HJB / free-boundary system in the viscosity sense.

Claim 2 (Lemma 2, Theorem 2). For the put-type exit payoff $g(x) = (K - x)_+$, a constant volatility and a non-increasing leader drift, the follower's exit region is a half-line. When the exit region is a half-line, the leader's equation satisfies a comparison principle: every sub-solution lies below and every super-solution above the leader's value.

Claim 3 (Theorem 3). In the alternating fixed-point algorithm each half-step is computed by sweeps that treat the jump term explicitly; each sweep is a sup-norm contraction with rate at most $\rho(q, \lambda) = \lambda / (q + \lambda)$. In the half-line case the outer iteration is a scalar map on the follower's threshold; it converges geometrically, to the unique equilibrium in an interval, whenever that map is a contraction there. Its rate depends on the payoffs, not only on q and λ .

Claim 4 (Proposition 1). For Kou double-exponential jumps the follower's best response to a constant drift is in closed form: the threshold is $b^* = K - 1/\beta_1 - 1/\beta_2 + 1/\eta_-$, where $-\beta_1, -\beta_2$ are the two negative roots of the Kou exponent equation, and the value is a sum of two exponentials. Section 8 computes both half-steps of the algorithm on a grid: the follower half-step against this closed form, and the leader half-step against the comparison theorem.

The paper is organised as follows. Section 2 fixes the probabilistic basis and the viscosity framework and states (A1)–(A3). Section 3 specifies the game and the coupled system. Sections 4 and 5 prove existence and verification under monotone stopping. Section 6 develops the alternating algorithm and its convergence. Section 7 specialises to Kou jumps, and Section 8 computes the two half-steps numerically. Section 9 concludes.



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 $dt \nu(dz)$. The state process $X_t \in \mathbb{R}$ is the unique strong solution of

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

where $\mu_L : \mathbb{R} \rightarrow \mathbb{R}$ is the leader's drift rule (specified in §3) and $\sigma : \mathbb{R} \rightarrow \mathbb{R}_+$ is a fixed Lipschitz diffusion coefficient. Because the jumps are compensated, μ_L is the mean rate of change of X .

2.2 Generator and viscosity solutions

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

$$\mathcal{L}f(x) = \mu_L(x) \partial_x f(x) + \frac{1}{2} \sigma^2(x) \partial_{xx} f(x) + \int_{\mathbb{R} \setminus \{0\}} [f(x+z) - f(x) - z \partial_x f(x)] \nu(dz). \quad (2.2)$$

The integral term is finite under (A2) below. We work with viscosity solutions in the sense of Crandall, Ishii and Lions [7], in the form adapted to non-local operators: a function u is a viscosity sub-solution of $F(x, u, \partial_x u, \partial_{xx} u, \mathcal{L}u) = 0$ if, for every C^2 test function φ such that $u - \varphi$ has a local maximum at x_0 , the inequality $F(x_0, u(x_0), \partial_x \varphi(x_0), \partial_{xx} \varphi(x_0), \mathcal{L}\varphi(x_0)) \leq 0$ holds; super-solution is the reverse inequality. Throughout the paper, "solution" without qualifier means "viscosity solution."

2.3 Function spaces

Two spaces are used. $C_{\text{lin}}(\mathbb{R})$ is the space of continuous functions of at most linear growth, with the weighted norm

$$\|u\|_w := \sup_{x \in \mathbb{R}} w(x) |u(x)|, \quad w(x) := (1 + |x|)^{-1}. \quad (2.3)$$

Differences of iterates in §6 are bounded, and for them we use the sup norm $\|u\|_\infty := \sup_x |u(x)|$. The value functions of the game are in $C_{\text{lin}}(\mathbb{R})$, because the payoffs of §7 grow linearly.



2.4 Standing assumptions (A1)–(A3)

Assumption 1 (standing assumptions on drift, jumps, and payoffs). The following three conditions hold throughout the paper.

(A1) *Drift admissibility.* The leader's drift rule $\mu_L : \mathbb{R} \rightarrow \mathbb{R}$ belongs to the admissible set

$$\mathcal{U}_L := \{\mu : \mathbb{R} \rightarrow [-\bar{\mu}, \bar{\mu}] : \mu \text{ is } L_\mu\text{-Lipschitz}\}, \quad (2.4)$$

for fixed constants $\bar{\mu}, L_\mu > 0$. The diffusion σ is L_σ -Lipschitz, bounded and uniformly elliptic: $\underline{\sigma} \leq \sigma(x) \leq \bar{\sigma}$ for all $x \in \mathbb{R}$, with $\underline{\sigma} > 0$.

(A2) *Jump measure.* The Lévy measure ν is finite and has a finite first moment:

$$\lambda := \nu(\mathbb{R} \setminus \{0\}) < \infty, \quad \int_{\mathbb{R} \setminus \{0\}} |z| \nu(dz) < \infty. \quad (2.5)$$

We write $\hat{\nu} := \nu/\lambda$ for the jump-size law and $m := \int z \nu(dz)$ for the compensator drift.

(A3) *Payoffs.* The follower's exit payoff $g : \mathbb{R} \rightarrow \mathbb{R}$ is non-increasing and 1-Lipschitz. The leader's running payoff $\pi_L(x, \mu)$ is continuous, Lipschitz in x uniformly in μ with constant L_π , of at most linear growth in x , and strictly concave in μ ; the leader's terminal payoff h_L is Lipschitz.

Remark 1. On (A3). The canonical exit payoff is the put $g(x) = (K - x)_+$ used in §7: the follower exits a falling state. It grows linearly as $x \rightarrow -\infty$, which is why the value functions are taken in $C_{\text{lin}}(\mathbb{R})$, as in [6]. A payoff that falls and then rises, such as $|x - x^*|$, typically produces an exit region with two components, one on each side, and lies outside the monotone-stopping analysis of §5 (Remark 5).

Under (A1)–(A2), the SDE (2.1) admits a unique strong solution X_t for each $\mu_L \in \mathcal{U}_L$, and X is a Feller process with càdlàg paths; see Protter [8].

3. Stackelberg game setup

3.1 Admissible strategies

The leader's admissible drift rules form the set $\mathcal{U}_L$ of (2.4). With the topology of uniform convergence on compacts, $\mathcal{U}_L$ is convex, and compact by the Arzelà–Ascoli theorem applied to the uniform bound $\bar{\mu}$ and the Lipschitz bound L_μ . The follower's admissible set $\mathcal{T}$ consists of all $(\mathcal{F}_t)$ -stopping times, with the convention $e^{-q\tau}g(X_\tau) := 0$ on $\{\tau = \infty\}$.

3.2 Value functions

Fix a discount rate $q > 0$. Given a leader drift rule $\mu \in \mathcal{U}_L$, the follower's value function is

$$V_F(x; \mu) := \sup_{\tau \in \mathcal{T}} \mathbb{E}_x[e^{-q\tau} g(X_\tau)], \quad (3.1)$$



where X solves (2.1) with drift μ , and the follower's exit region is $\mathcal{S}(\mu) := \{x : V_F(x; \mu) \leq g(x)\}$. For a closed exit region $\mathcal{S}$ write $\tau_{\mathcal{S}} := \inf\{t \geq 0 : X_t \in \mathcal{S}\}$. The leader's payoff from a rule μ against the exit region $\mathcal{S}$, and the leader's value against $\mathcal{S}$, are

$$\begin{aligned} J_L(x; \mu, \mathcal{S}) &:= \mathbb{E}_x \left[\int_0^{\tau_{\mathcal{S}}} e^{-qs} \pi_L(X_s, \mu(X_s)) ds + e^{-q\tau_{\mathcal{S}}} h_L(X_{\tau_{\mathcal{S}}}) \right], \\ V_L(x; \mathcal{S}) &:= \sup_{\mu \in \mathcal{U}_L} J_L(x; \mu, \mathcal{S}), \end{aligned} \quad (3.2)$$

with the terminal term read as 0 on $\{\tau_{\mathcal{S}} = \infty\}$.

3.3 Equilibrium concept

Definition 1 (feedback Stackelberg equilibrium). A pair $(\mu_L^*, \mathcal{S}^*) \in \mathcal{U}_L \times \{\text{closed subsets of } \mathbb{R}\}$ is a *feedback Stackelberg equilibrium* if: *Step 1.* The follower best-responds to the announced rule: $\mathcal{S}^* = \mathcal{S}(\mu_L^*)$, so that $\tau^* = \tau_{\mathcal{S}^*}$ is optimal in (3.1) for the drift μ_L^* . *Step 2.* The leader's rule is optimal against the follower's exit rule: $J_L(x; \mu_L^*, \mathcal{S}^*) = V_L(x; \mathcal{S}^*)$ for every $x \in \mathbb{R}$. The equilibrium value functions are $V_F^* := V_F(\cdot; \mu_L^*)$ and $V_L^* := V_L(\cdot; \mathcal{S}^*)$.

Remark 2 (feedback and commitment). The leader moves first in the sense that the follower observes the rule μ_L^* and responds to it; Step 2 asks the leader's rule to be optimal against the follower's exit rule $\mathcal{S}^*$, not against the reaction map $\mu \mapsto \mathcal{S}(\mu)$. This is the feedback notion, and it is the one the coupled system (3.3) characterises. The commitment problem $\sup_{\mu} J_L(x; \mu, \mathcal{S}(\mu))$, in which the leader optimises through the follower's reaction map, is a bilevel problem; since μ_L^* is admissible in it, its value is at least $V_L^*(x)$. It is not treated here.

3.4 The coupled HJB / free-boundary system

Applying the dynamic-programming principle to (3.2) and the obstacle formulation of optimal stopping to (3.1), the equilibrium value functions satisfy the coupled system formalised next.

Definition 2 (coupled HJB / free-boundary system). A pair of continuous functions (V_L^*, V_F^*) of at most linear growth, together with a leader drift $\mu_L^* \in \mathcal{U}_L$ and a closed exit region $\mathcal{S}^* \subset \mathbb{R}$, solves the *coupled HJB / free-boundary system* if equations (3.3) below hold, where $\mathcal{L}^\mu$ denotes the generator (2.2) with the drift fixed at the value μ , and $\mathcal{S}^* = \{x : V_F^*(x) = g(x)\}$.

The coupled system is



$$\begin{aligned}
\sup_{\mu \in [-\bar{\mu}, \bar{\mu}]} [\mathcal{L}^\mu V_L^*(x) + \pi_L(x, \mu)] - qV_L^*(x) &= 0, & x \in \mathbb{R} \setminus \mathcal{S}^*, \\
V_L^*(x) &= h_L(x), & x \in \mathcal{S}^*, \\
\min[qV_F^*(x) - \mathcal{L}^{\mu_L^*} V_F^*(x), V_F^*(x) - g(x)] &= 0, & x \in \mathbb{R},
\end{aligned} \tag{3.3}$$

with $\mu_L^*(x)$ attaining the supremum in the first line. The system couples the leader's HJB equation (first two lines) to the follower's variational inequality (third line) through the dependence of $\mathcal{S}^*$ on V_F^* and of $\mathcal{L}^{\mu_L^*} V_F^*$ on μ_L^* . The leader's equation is posed on the open set $\mathbb{R} \setminus \mathcal{S}^*$, but the non-local term of $\mathcal{L}^\mu$ reads V_L^* on all of $\mathbb{R}$, including the values h_L on $\mathcal{S}^*$ that a jump can reach.

Remark 3. When $\mathcal{S}^*$ is a half-line (the monotone-stopping case of §5), the first two lines of (3.3) are an HJB equation on a half-line with the data h_L prescribed on the complementary half-line. This is the structure the verification theorem of §5 and the threshold iteration of §6 use.

Figure 1 summarises the coupling: the leader's value is computed against the follower's exit region $\mathcal{S}^*$, which is determined by the follower's value V_F^* , which depends on the leader's drift rule μ_L^* , which is the maximiser in the leader's HJB equation. The existence proof of §4 analyses this fixed-point structure directly.




Figure 1: The structure of the game. The leader announces a drift rule for a jump-diffusion state; the follower observes the state and the rule and exits on entering its exit region. Each rule is a best response to the other (the arc), which makes the equilibrium a joint fixed point of the two best-response maps.

4. Existence of an equilibrium

This section proves Claim 1 of §1. The proof constructs a coupling map on $\mathcal{U}_L$, establishes its continuity under a hypothesis on the best responses, and applies the Schauder–



Tychonoff fixed-point theorem.

4.1 The coupling map Φ

Define the **follower best response** $\Phi_F(\mu) := \mathcal{S}(\mu)$, the exit region of (3.1) under the drift μ , and the **leader best response** to a closed exit region $\mathcal{S}$,

$$\Phi_L(\mathcal{S})(x) := \arg \max_{\mu \in [-\bar{\mu}, \bar{\mu}]} [\mu \partial_x V_L(x; \mathcal{S}) + \pi_L(x, \mu)], \quad x \notin \mathcal{S}, \quad (4.1)$$

since the drift enters the generator only through the term $\mu \partial_x$. Strict concavity of π_L in μ (A3) makes the maximiser unique. On $\mathcal{S}$ the drift affects neither value, and $\Phi_L(\mathcal{S})$ is extended there by its one-sided limits at $\partial \mathcal{S}$. The composition $\Phi := \Phi_L \circ \Phi_F$ maps drift rules to drift rules, and its fixed points are the leader rules of feedback Stackelberg equilibria.

Hypothesis (H). (H1) The map $\mu \mapsto \mathcal{S}(\mu)$ is continuous from $\mathcal{U}_L$, with the topology of uniform convergence on compacts, into the closed sets with the Hausdorff metric on compacts. (H2) Φ maps $\mathcal{U}_L$ into itself: every leader best response $\Phi(\mu)$, $\mu \in \mathcal{U}_L$, is L_μ -Lipschitz.

(H1) excludes a degenerate contact $V_F - g$ at the free boundary, where the exit region can jump as the drift moves; in the half-line case it is continuity of the threshold in the drift. (H2) is a regularity statement on $\partial_x V_L$ on the closure of the continuation region, uniform over exit regions; the constant L_μ in (A1) has to be chosen at least as large as the resulting Lipschitz constant.

4.2 Continuity and compactness

Lemma 1 (Continuity of the coupling map). Under (A1)–(A3) and (H), the map $\Phi : \mathcal{U}_L \rightarrow \mathcal{U}_L$ is continuous in the topology of uniform convergence on compacts.

Proof. Sketch. Step 1. (Continuity of Φ_F .) Let $\mu^{(n)} \rightarrow \mu^\infty$ uniformly on compacts. By the stability of strong solutions of SDEs with Lipschitz coefficients [8], $X_t^{\mu^{(n)}} \rightarrow X_t^{\mu^\infty}$ in probability uniformly on $[0, T]$ for every $T > 0$. Since g is Lipschitz and the discount $e^{-q\tau}$ controls the linear growth of g , dominated convergence gives $V_F(\cdot; \mu^{(n)}) \rightarrow V_F(\cdot; \mu^\infty)$ uniformly on compacts, and (H1) gives $\Phi_F(\mu^{(n)}) \rightarrow \Phi_F(\mu^\infty)$.

Step 2. (Continuity of Φ_L .) For exit regions converging in the Hausdorff metric on compacts, the leader values $V_L(\cdot; \mathcal{S})$ and their derivatives converge locally uniformly on compact subsets of the continuation regions, by the stability of viscosity solutions and the uniform ellipticity of (A1) [2]. The maximiser in (4.1) is a continuous function of $\partial_x V_L$ because π_L is strictly concave in μ . Hence $\Phi = \Phi_L \circ \Phi_F$ is continuous. $\square$



4.3 Existence by Schauder–Tychonoff

Theorem 1 (Existence of an equilibrium). Under **(A1)–(A3)** and **(H)**, there is a feedback Stackelberg equilibrium $(\mu_L^*, \mathcal{S}^*)$. Its value functions (V_L^*, V_F^*) are continuous, of at most linear growth, and solve the coupled system (3.3) in the viscosity sense.

Proof. Sketch. Step 1. (Fixed point.) $\mathcal{U}_L$ is a non-empty, convex, compact subset of the locally convex space $C(\mathbb{R})$ with the topology of uniform convergence on compacts (§3), and Φ maps it continuously into itself by (H2) and Lemma 1. The Schauder fixed-point theorem [9], in Tychonoff’s form for locally convex spaces, gives $\mu_L^* \in \mathcal{U}_L$ with $\Phi(\mu_L^*) = \mu_L^*$.

Step 2. (Equilibrium.) Set $\mathcal{S}^* := \Phi_F(\mu_L^*)$; Step 1 of Definition 1 holds by construction. By the dynamic-programming principle [2], $V_L(\cdot; \mathcal{S}^*)$ is a viscosity solution of the first two lines of (3.3), and the feedback maximiser $\mu_L^* = \Phi_L(\mathcal{S}^*)$, which is Lipschitz by (H2), is an optimal Markov control for it; this is Step 2 of Definition 1.

Step 3. (The follower’s equation.) $V_F^* = V_F(\cdot; \mu_L^*)$ solves the third line of (3.3) by the dynamic-programming principle for optimal stopping; for a constant drift this is Theorem 6 of [6], and the argument is the same for a Lipschitz drift. Both values have at most linear growth because the payoffs do (A3) and $\mathbb{E}_x[\sup_t e^{-qt}|X_t|] < \infty$ under (A1)–(A2). $\square$

Remark 4 (uniqueness). Theorem 1 does not assert uniqueness. In the half-line case of §5 the equilibria are the fixed points of a scalar threshold map, and Theorem 3 (ii) gives uniqueness in any interval on which that map is a contraction.

5. Verification under monotone stopping

Theorem 1 gives existence in the viscosity sense but leaves open how a candidate leader value can be checked. This section closes that gap when the follower’s exit region is a half-line, and gives a sufficient condition for that structure.

5.1 Half-line exit regions

Lemma 2 (Half-line exit region). Let $\sigma \equiv \sigma_0 > 0$ be constant, let $g(x) = (K - x)_+$, and let $\mu \in \mathcal{U}_L$ be non-increasing (for example constant). Then the follower’s exit region $\mathcal{S}(\mu)$ is empty or a closed half-line $(-\infty, b]$ with $b < K$.

Proof. For $y < x$ drive X^x and X^y by the same Brownian motion and the same Poisson random measure. With constant σ and additive jumps, $D_t := X_t^x - X_t^y$ solves the ordinary differential equation $\dot{D}_t = \mu(X_t^x) - \mu(X_t^y)$ with $D_0 = x - y > 0$. Since μ is Lipschitz, D cannot change sign, and since μ is non-increasing, $\dot{D}_t \leq 0$ while $D_t \geq 0$; hence $0 \leq D_t \leq x - y$ for all t . For any stopping time τ , $X_\tau^y \geq X_\tau^x - (x - y)$, and because g is non-increasing and 1-Lipschitz, $g(X_\tau^y) \leq g(X_\tau^x) + (x - y)$; taking expectations and the supremum over τ gives $V_F(y) \leq V_F(x) + (x - y)$. If $x \in \mathcal{S}(\mu)$ and $x \leq K$, this gives $V_F(y) \leq K - y = g(y)$, so $y \in \mathcal{S}(\mu)$. No $x \geq K$ lies in $\mathcal{S}(\mu)$, since $g(x) = 0 < V_F(x)$: with $\sigma_0 > 0$ the state enters $(-\infty, K)$ before $\mathbf{e}(q)$ with positive probability. $\mathcal{S}(\mu)$ is closed because V_F and g are continuous. For a constant drift this is Proposition 1 of [6]. $\square$



Remark 5 (when monotone stopping fails). The proof uses the shape of the put payoff (slope -1 wherever it is positive) and the order-preserving flow of a non-increasing drift. For a payoff that falls and then rises the exit region typically has two components, and for an increasing drift the flow can spread nearby states apart. Theorem 2 and the threshold iteration of §6 therefore take the half-line structure as a hypothesis, which Lemma 2 guarantees in the cases above. Two-sided exit regions need a different technique, such as the quasi-variational inequalities of Bensoussan and Lions [10].

5.2 Comparison principle for the leader's HJB equation

Theorem 2 (Viscosity verification under monotone stopping). Let $\mathcal{S}^* = (-\infty, b^*]$ be a half-line, and let $\underline{V}, \bar{V} \in C_{\text{lin}}(\mathbb{R})$ be, respectively, a viscosity sub-solution and a viscosity super-solution of

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

with $\underline{V} \leq h_L \leq \bar{V}$ on $\mathcal{S}^*$. Then $\underline{V} \leq V_L(\cdot; \mathcal{S}^*) \leq \bar{V}$ on $\mathbb{R}$. In particular (5.1) with the data h_L on $\mathcal{S}^*$ has exactly one solution in $C_{\text{lin}}(\mathbb{R})$, namely $V_L(\cdot; \mathcal{S}^*)$.

Proof. Sketch. By the dynamic-programming principle $V_L(\cdot; \mathcal{S}^*)$ is itself a solution with the data h_L on $\mathcal{S}^*$, so it suffices to show $u \leq v$ for a sub-solution u and a super-solution v with $u \leq v$ on $\mathcal{S}^*$.

Step 1. (Doubling.) Suppose $M := \sup_x (u - v)(x) > 0$. For $\varepsilon, \delta > 0$ let $(x_\varepsilon, y_\varepsilon)$ maximise $\Psi(x, y) := u(x) - v(y) - (x - y)^2 / (2\varepsilon) - \delta(x^2 + y^2)$; the quadratic penalty makes a maximiser exist for functions of linear growth. For δ small the maximum exceeds $M/2$, and since $u \leq v$ on $\mathcal{S}^*$ and both functions are continuous at the single boundary point b^* , the points $x_\varepsilon, y_\varepsilon$ lie in the open half-line (b^*, ∞) for ε small.

Step 2. (Doubling estimate.) The non-local form of the Crandall–Ishii lemma, used for the same purpose in Theorem 5 of [6], gives, with the same control value μ in both inequalities,

$$q(u(x_\varepsilon) - v(y_\varepsilon)) \leq \frac{3}{2\varepsilon} [\sigma(x_\varepsilon) - \sigma(y_\varepsilon)]^2 + L_\pi |x_\varepsilon - y_\varepsilon| + R_\varepsilon + O(\delta), \quad (5.2)$$

where the drift terms cancel because the control value does not depend on the state, L_π comes from the Lipschitz bound on π_L , and R_ε is the jump residual, bounded by $\int |z| \nu(dz)$ times a quantity that vanishes with $|x_\varepsilon - y_\varepsilon|^2 / \varepsilon$, using (A2).

Step 3. (Limit.) By (A1), $[\sigma(x_\varepsilon) - \sigma(y_\varepsilon)]^2 \leq L_\sigma^2 |x_\varepsilon - y_\varepsilon|^2$, and the standard estimate $|x_\varepsilon - y_\varepsilon|^2 / \varepsilon \rightarrow 0$ makes the right-hand side of (5.2) vanish as $\varepsilon \downarrow 0$ and then $\delta \downarrow 0$, while the left-hand side stays above $qM/2$. This contradiction gives $u \leq v$. $\square$

The half-line hypothesis enters Step 1: the boundary of the continuation region is the single point b^* , where the data h_L are attained continuously because $\sigma \geq \underline{\sigma} > 0$. Theorem 2 is what identifies the output of the leader half-step of §6 with $V_L(\cdot; \mathcal{S})$.



6. Fixed-point algorithm and convergence

The verification theorem of §5 fixes the analytic structure of the equilibrium but does not compute it. This section states the alternating fixed-point iteration, splits off the jump term to compute each half-step, and proves what converges at which rate.

6.1 The alternating iteration

Assume that every exit region produced below is a half-line (Lemma 2 or Remark 5). Start from a threshold b_0 and, for $n = 0, 1, 2, \dots$:

Leader half-step. With $\mathcal{S}_n := (-\infty, b_n]$, compute $V_L^{(n)} := V_L(\cdot; \mathcal{S}_n)$, the solution of (5.1) with the data h_L on $\mathcal{S}_n$ (Theorem 2), and the drift rule $\mu_n := \Phi_L(\mathcal{S}_n)$ of (4.1).

Follower half-step. Compute $V_F^{(n+1)} := V_F(\cdot; \mu_n)$, the solution of the obstacle problem in the third line of (3.3) under the drift μ_n , and its threshold b_{n+1} , with $\mathcal{S}(\mu_n) = (-\infty, b_{n+1}]$.

This defines the threshold map $\Gamma : b_n \mapsto b_{n+1}$. By Definition 1, the equilibria with a half-line exit region are exactly the pairs $(\Phi_L((-\infty, b^*]), (-\infty, b^*])$ with $\Gamma(b^*) = b^*$.

6.2 Splitting off the jumps

Under (A2) the generator (2.2) splits into a diffusion part and a jump part,

$$\begin{aligned} \mathcal{L}^\mu f &= \mathcal{D}^\mu f + \lambda(\hat{K}f - f), \\ \mathcal{D}^\mu f &:= (\mu - m) \partial_x f + \frac{1}{2} \sigma^2 \partial_{xx} f, \quad \hat{K}f(x) := \int f(x+z) \hat{\nu}(dz), \end{aligned} \quad (6.1)$$

where $\hat{K}$ is a Markov operator: $\hat{K}f \geq 0$ for $f \geq 0$ and $\hat{K}1 = 1$. Write Y for the diffusion with generator $\mathcal{D}^\mu$, without jumps. Each half-step is computed by sweeps in which the jump term is evaluated at the previous sweep. The **follower sweep** under a drift μ maps u to the solution $w = F_\mu(u)$ of the diffusion obstacle problem

$$\begin{aligned} \min[(q + \lambda)w - \mathcal{D}^\mu w - \lambda \hat{K}u, w - g] &= 0, \quad \text{equivalently} \\ F_\mu(u)(x) &= \sup_\tau \mathbb{E}_x \left[\int_0^\tau e^{-(q+\lambda)t} \lambda \hat{K}u(Y_t) dt + e^{-(q+\lambda)\tau} g(Y_\tau) \right]. \end{aligned} \quad (6.2)$$

The **leader sweep** against a half-line $\mathcal{S}$ maps u , extended by h_L on $\mathcal{S}$, to the solution $w = G_{\mathcal{S}}(u)$ of

$$\sup_{\mu \in [-\bar{\mu}, \bar{\mu}]} [\mathcal{D}^\mu w + \pi_L(\cdot, \mu)] + \lambda \hat{K}u - (q + \lambda)w = 0 \quad \text{on } \mathbb{R} \setminus \mathcal{S}, \quad w = h_L \quad \text{on } \mathcal{S}. \quad (6.3)$$

In (6.2) the jump term enters as a running reward at the rate λ and the discount rises from q to $q + \lambda$: up to the first jump the state moves as Y , and at the first jump, an



exponential time of rate λ , the follower receives the value u at the post-jump state.

6.3 Convergence

Theorem 3 (Convergence of the alternating iteration). Assume **(A1)**–**(A3)**.

(i) (*Half-steps.*) For every drift rule $\mu \in \mathcal{U}_L$, every half-line $\mathcal{S}$, and all u, v with $u - v$ bounded,

$$\|F_\mu(u) - F_\mu(v)\|_\infty \leq \rho \|u - v\|_\infty, \quad \|G_{\mathcal{S}}(u) - G_{\mathcal{S}}(v)\|_\infty \leq \rho \|u - v\|_\infty, \quad (6.4)$$

with $\rho(q, \lambda) := \lambda/(q + \lambda) \in (0, 1)$. The unique fixed points are $V_F(\cdot; \mu)$ and $V_L(\cdot; \mathcal{S})$, and the sweeps converge to them geometrically: after k sweeps the sup-norm error is at most ρ^k times the initial one.

(ii) (*Outer iteration.*) If Γ maps a closed interval I into itself and $|\Gamma(b) - \Gamma(b')| \leq \rho_\Gamma |b - b'|$ on I with $\rho_\Gamma < 1$, then Γ has exactly one fixed point b^* in I , the outer iteration started anywhere in I satisfies $|b_n - b^*| \leq \rho_\Gamma^n |b_0 - b^*|$, and $(\Phi_L((-\infty, b^*]), (-\infty, b^*))$ is the only equilibrium with a half-line exit region whose threshold lies in I .

Proof. Step 1. (Follower sweep.) By the representation in (6.2) and $|\sup_\tau a_\tau - \sup_\tau b_\tau| \leq \sup_\tau |a_\tau - b_\tau|$,

$$\begin{aligned} |F_\mu(u)(x) - F_\mu(v)(x)| &\leq \sup_\tau \mathbb{E}_x \left[\int_0^\tau e^{-(q+\lambda)t} \lambda |\hat{K}(u - v)(Y_t)| dt \right] \\ &\leq \lambda \|u - v\|_\infty \int_0^\infty e^{-(q+\lambda)t} dt = \rho \|u - v\|_\infty, \end{aligned}$$

because $\hat{K}$ is a Markov operator, so $\|\hat{K}f\|_\infty \leq \|f\|_\infty$.

Step 2. (Leader sweep.) $G_{\mathcal{S}}(u)(x) = \sup_\mu \mathbb{E}_x [\int_0^{\tau_{\mathcal{S}}} e^{-(q+\lambda)t} (\pi_L(Y_t, \mu(Y_t)) + \lambda \hat{K}u(Y_t)) dt + e^{-(q+\lambda)\tau_{\mathcal{S}}} h_L(Y_{\tau_{\mathcal{S}}})]$, with $\tau_{\mathcal{S}}$ the entrance time of Y into $\mathcal{S}$ (the comparison argument of Theorem 2, applied to the diffusion Y , identifies the solution of (6.3) with this value). The payoffs π_L and h_L cancel in the difference, and the estimate of Step 1 applies with the supremum over drift rules in place of the supremum over stopping times. Jumps that land in $\mathcal{S}$ are accounted for by the extension $u = h_L$ on $\mathcal{S}$.

Step 3. (Fixed points.) Conditioning on the first jump time, an independent exponential time of rate λ up to which X moves as Y , the dynamic-programming principle gives $F_\mu(V_F(\cdot; \mu)) = V_F(\cdot; \mu)$ and $G_{\mathcal{S}}(V_L(\cdot; \mathcal{S})) = V_L(\cdot; \mathcal{S})$. The differences between an iterate and these fixed points are bounded when the sweeps start from g (respectively from any function equal to h_L on $\mathcal{S}$ with the growth of V_L), so the Banach fixed-point theorem on the affine space of such functions gives uniqueness and the geometric rate.

Step 4. (Outer iteration.) This is the Banach fixed-point theorem for the map Γ on the complete metric space I , combined with the characterisation of equilibria as fixed points of Γ . $\square$

Remark 6 (rate degeneracy as $q \downarrow 0$). The rate $\rho(q, \lambda) = \lambda/(q + \lambda)$ approaches 1 as $q \downarrow 0$.



The undiscounted game is therefore not covered by Theorem 3 and requires a separate ergodic-control treatment.

Remark 7 (which rate belongs where). As $\lambda \downarrow 0$ the rate ρ tends to 0: without jumps a single sweep is the exact solve of the half-step. The rate $\lambda/(q + \lambda)$ therefore belongs to the half-steps and cannot be a rate for the outer iteration, since at $\lambda = 0$ it would assert that one outer step reaches the equilibrium from any initial threshold. The outer rate ρ_Γ is the product of two sensitivities, of the leader's best-response drift to the threshold and of the follower's threshold to the drift, and it depends on the payoffs. When $\rho_\Gamma \geq 1$ the plain iteration can fail although an equilibrium exists; the equilibrium threshold is then a root of the scalar function $\Gamma(b) - b$, and a bracketing search replaces the iteration.

6.4 Algorithmic implementation

Algorithm 1 (Appendix A) implements the iteration on a uniform grid: central differences for $\mathcal{D}^\mu$; the jump term $\hat{K}u$ of a piecewise-linear iterate integrated exactly against the jump law (Appendix B); each diffusion obstacle problem (6.2) solved exactly on the grid by Howard's policy iteration [11] on a tridiagonal system; and the leader sweep (6.3) by policy iteration over the maximiser. A sweep costs $O(N)$ operations per policy-iteration step for N grid points. By Theorem 3 (i) a half-step needs about $\log(1/\eta)/\log(1/\rho)$ sweeps to reach a tolerance η ; the number of outer iterations is governed by ρ_Γ , which is not known in advance. Section 8 computes the follower half-step.

7. Worked example: Kou jumps and a put-type exit

This section specialises the framework to the **Kou double-exponential jump model** [12], with the put as the follower's exit payoff. For a constant leader drift the follower's best response is then in closed form, through the Wiener–Hopf factorisation [13] in the form proved for Kou processes in [6].

7.1 Kou specialisation

Specialise (2.1) by setting

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

with up-rate $\eta_+ > 0$, down-rate $\eta_- > 0$, up-jump probability $p \in (0, 1)$, and total intensity $\lambda > 0$. The measure (7.1) satisfies **(A2)**, with $\hat{\nu}$ the two-sided exponential law and compensator drift $m = \lambda(p/\eta_+ - (1 - p)/\eta_-)$. Specialise the payoffs to

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

with $\alpha, \beta \in \mathbb{R}$, $\gamma > 0$, $K > 0$, $\kappa \geq 0$. They satisfy **(A3)**, and the leader's best response (4.1) is the clipped first-order condition



$$\Phi_L(\mathcal{S})(x) = \min\left\{\bar{\mu}, \max\left\{-\bar{\mu}, \frac{\beta + \partial_x V_L(x; \mathcal{S})}{\gamma}\right\}\right\}. \quad (7.3)$$

7.2 The Kou exponent and the law of the infimum

For a constant drift μ the state is a Lévy process. The compensated jumps of (2.1) contribute $-m$ to its drift, so its Laplace exponent is the Kou exponent of [6] with drift $b := \mu - m$:

$$\psi(r) := \log \mathbb{E}_0[e^{rX_1}] = br + \frac{1}{2}\sigma_0^2 r^2 + \lambda \left(\frac{p\eta_+}{\eta_+ - r} + \frac{(1-p)\eta_-}{\eta_- + r} - 1 \right), \quad -\eta_- < r < \eta_+. \quad (7.4)$$

Lemma 3 (Kou roots and the law of the infimum). Let $q > 0$ and μ be constant. The equation $\psi(r) = q$, continued as a rational function, has exactly four real roots $-\beta_2 < -\eta_- < -\beta_1 < 0 < \alpha_1 < \eta_+ < \alpha_2$. The infimum $I := \inf_{0 \leq s \leq e(q)} X_s - X_0$ has the Wiener–Hopf factor

$$\Phi_q^-(r) := \mathbb{E}[e^{rI}] = \frac{\beta_1 \beta_2 (\eta_- + r)}{\eta_- (\beta_1 + r)(\beta_2 + r)}, \quad r > -\beta_1, \quad (7.5)$$

so that $-I$ is the mixture $w_1 \text{Exp}(\beta_1) + w_2 \text{Exp}(\beta_2)$ with $w_1 = \beta_2(\eta_- - \beta_1)/(\eta_-(\beta_2 - \beta_1))$ and $w_2 = \beta_1(\beta_2 - \eta_-)/(\eta_-(\beta_2 - \beta_1))$, both positive, and $\mathbb{E}[-I] = 1/\beta_1 + 1/\beta_2 - 1/\eta_-$.

Proof. The roots and the factor are Theorem 2 of [6]; the identity $q/(q - \psi(r)) = \Phi_q^+(r)\Phi_q^-(r)$ behind them is its Theorem 1. The partial-fraction form of (7.5), $\sum_i w_i \beta_i / (\beta_i + r)$, gives the mixture, with $w_1 + w_2 = 1$; the mean is $-\partial_r \log \Phi_q^-(0)$. $\square$

7.3 The follower's best response to a constant drift

Proposition 1 (The follower's best response under Kou jumps). For the Kou specialisation (7.1) with $g(x) = (K - x)_+$ and a constant leader drift $\mu_L \equiv \mu$, the follower's exit region is $\mathcal{S}(\mu) = (-\infty, b^*]$ with

$$b^* = K - \frac{1}{\beta_1} - \frac{1}{\beta_2} + \frac{1}{\eta_-}, \quad V_F(x; \mu) = \begin{cases} K - x, & x \leq b^*, \\ \frac{w_1}{\beta_1} e^{-\beta_1(x-b^*)} + \frac{w_2}{\beta_2} e^{-\beta_2(x-b^*)}, & x > b^*, \end{cases} \quad (7.6)$$

where $\beta_1, \beta_2, w_1, w_2$ are those of Lemma 3 for the drift $b = \mu - m$. The value is continuously differentiable across b^* (smooth fit).

Proof. For a constant drift X is a Lévy process that is not a subordinator, and Theorem 7 of [6] gives $b^* = K + \mathbb{E}[I]$ and $V_F(x) = \mathbb{E}[(b^* - x - I)^+]$ for the put (the Novikov–Shiryayev form). Lemma 3 gives $\mathbb{E}[I]$, which is the threshold, and for $c := x - b^* > 0$ and $M \sim \text{Exp}(\beta)$, $\mathbb{E}[(M - c)^+] = e^{-\beta c}/\beta$, which gives the value on $x > b^*$. Its derivative at b^{*+} is $-(w_1 + w_2) = -1$, the slope of $K - x$. $\square$



Remark 8 (what the closed form gives). Proposition 1 gives the follower's best response to every constant drift: the follower half-step of §6 started from a constant rule, and the follower's reaction map $\mu \mapsto b^*(\mu)$ when the leader is restricted to constant rules. It does not give the equilibrium of the general game. The leader's best response (7.3) is constant only if $\partial_x V_L$ is constant on the continuation region, which the data at b^* generically rule out; under a state-dependent drift X is not a Lévy process and the Wiener–Hopf closed forms do not apply. The equilibrium is then computed by the algorithm of §6, with Proposition 1 as an exact check of its follower half-step.

8. Numerical illustration: the two half-steps

This section computes the two half-steps of §6 as Algorithm 1 states them: the follower half-step at a constant leader drift, checked against Proposition 1, and the leader half-step against the exit region that the follower half-step produces, checked against Theorem 2. Both test the rate bound of Theorem 3 (i). The outer threshold iteration is not run here. The computation is `python/follower_check.py`; the figures are drawn by `python/-make_figure.py`.

8.1 Set-up

The Kou parameters are the benchmark of [6]: $\sigma_0 = 0.20$, $p = 0.6$, $\eta_+ = 10$, $\eta_- = 5$, with $K = 1$ and $q = 0.05$; they are not calibrated to data. The jump intensity takes the values $\lambda \in \{0.5, 1, 2\}$, and the leader drift is $\mu_L \equiv 0$ (no intervention: the state is a martingale), so that the drift of the Lévy exponent is $b = -m$.

Method 1 is Proposition 1. The roots $\beta_1 \in (0, \eta_-)$ and $\beta_2 \in (\eta_-, \infty)$ are found by bracketing; as a check, the four eigenvalues of the companion matrix of the quartic numerator of $\psi(r) - q$ agree with the bracketed roots to within 5.3×10^{-15} . At the drift $b = 0$ and $\lambda = 1$, Method 1 returns $b^* = -0.0897$ with $(\beta_1, \beta_2) = (0.8563, 8.2010)$, the values reported in [6].

Method 2 is Algorithm 1's follower half-step: the sweeps (6.2) on the grid $[-8, 40]$ with step $h \in \{0.01, 0.005\}$, with $u = g$ below the grid and $u = 0$ above it, started from $u_0 = g$ and stopped when the sup-norm change $d_k := \|u_k - u_{k-1}\|_\infty$ falls below 10^{-12} . It knows the generator and the payoff, not the roots or the closed form. Each sweep's obstacle problem is solved exactly on the grid by policy iteration.



Table 1: Method 1 at the benchmark with the constant leader drift $\mu_L \equiv 0$: the drift of the Lévy exponent, the two negative-root rates of Lemma 3, the first mixture weight of $-I$, and the follower's threshold of Proposition 1.

λ	drift b	β_1	β_2	w_1	b^*
0.5	+0.010	1.2259	6.9295	0.9171	0.2400
1	+0.020	1.0464	8.5476	0.9010	0.1273
2	+0.040	0.8476	11.3406	0.8976	-0.0680

8.2 Results

Table 2: Method 2, the contraction of the follower sweeps at $h = 0.005$. The ratios are those of successive sup-norm changes, d_{k+1}/d_k : the largest over the sweeps with $d_k > 10^{-10}$, and the asymptotic ratio as the geometric mean over the sweeps with $10^{-10} < d_k < 10^{-4}$. "Sweeps by the bound" is the smallest k with $d_1 \rho^{k-1} < 10^{-12}$.

λ	$\rho = \lambda/(q + \lambda)$	largest ratio	asymptotic ratio	sweeps	sweeps by the bound
0.5	0.9091	0.9040	0.8995	216	268
1	0.9524	0.9496	0.9469	411	520
2	0.9756	0.9741	0.9726	793	1022

Table 3: Method 2 against Method 1: the largest value error $\max|V_h - V_F|$ over the grid points $x \leq 10$ at two step sizes and their ratio, and the grid threshold b_h , the vertex of the parabola through $V_h - g$ at the first three continuation nodes, against b^* .

λ	error, $h = 0.01$	error, $h = 0.005$	error ratio	$ b_h - b^* $, $h = 0.005$
0.5	2.0×10^{-5}	5.1×10^{-6}	4.00	1.6×10^{-4}
1	3.1×10^{-5}	7.2×10^{-6}	4.34	1.5×10^{-4}
2	5.1×10^{-5}	1.2×10^{-5}	4.16	2.6×10^{-4}

Three findings. First, every ratio of successive changes is below ρ , as Theorem 3 (i) requires, and the asymptotic ratio lies just below the bound, so the bound is close to sharp for the follower half-step: the half-step needs 216, 411 and 793 sweeps to reach 10^{-12} for $\lambda = 0.5, 1, 2$, the same at both step sizes, against 268, 520 and 1022 predicted by the bound. Second, the grid value converges to the closed form at second order: halving



h divides the error by 4.00, 4.34 and 4.16. Third, the grid threshold agrees with b^* to within 2.6×10^{-4} at $h = 0.005$. The grid solves the third line of the coupled system (3.3) under the drift $\mu_L \equiv 0$, and its agreement with the stopping value of Proposition 1 is the follower half of the statement of Theorem 1 that the value functions solve the coupled system. At every step size the computed exit region is a single interval ending at b_h , as Lemma 2 requires. Figure 2 shows both comparisons. At this drift, more jumps lower the threshold, so the follower waits longer before exiting, and slow the sweeps, as the bound $\rho = \lambda/(q + \lambda)$ predicts.




Figure 2: The follower half-step of Algorithm 1 against Proposition 1, Kou benchmark, constant leader drift $\mu_L \equiv 0$. Left: the follower's value V_F from the closed form (7.6) for $\lambda = 0.5, 1, 2$ (lines; triangles at the thresholds $b^* = 0.240, 0.127, -0.068$) and the converged grid sweeps at $h = 0.01$ (circles). Right: the sup-norm change per sweep d_k (lines) against the bound $d_1 \rho^{k-1}$ of Theorem 3 (i) (thin grey lines in the same dash pattern).

8.3 The leader half-step

For Theorem 2 we compute the leader half-step against the exit region $(-\infty, b^*]$ found above for $\lambda = 1$, $b^* = 0.1273$, with the payoffs (7.2) at $\alpha = 0.1$, $\beta = 0.5$, $\gamma = 1$, $\kappa = 0$ and the drift bound $\bar{\mu} = 3$. Each sweep (6.3) is solved on the grid by policy iteration over the clipped first-order condition (7.3), which settles in at most 9 steps; the data at the top of the grid are the exit-free value $\alpha x/q + C$ with $C = (\alpha/q + \beta)^2/(2\gamma q) = 62.5$.

Started from the exit-free value and from zero, the sweeps reach a change below 10^{-10} after 12 and 53 sweeps and end within 4.8×10^{-11} of each other, as the uniqueness part of Theorem 2 requires. The largest ratio of successive changes over both starts is 0.912, below $\rho = 0.952$, which is Theorem 3 (i) for the leader sweep. The payoffs J_c of the constant drifts $c \in \{0, 2.5, 3\}$, computed by the same sweeps, are sub-solutions with the same data on the exit region, and Theorem 2 puts them below the value: the smallest margins $V_L - J_c$ over the continuation nodes are 33.5, -1.9×10^{-11} and 1.32 at $h = 0.005$. The second is zero to within the sweep tolerance, because far from the exit the best response tends to $(\alpha/q + \beta)/\gamma = 2.5$.



The leader's best-response drift is not constant: on the continuation nodes up to $x = 10$ it falls from 3.00, the bound, near the exit to 2.50 far from it. The next follower half-step therefore faces a state-dependent drift, and Proposition 1 no longer applies (Remark 8); since this drift is non-increasing, Lemma 2 still makes the next exit region a half-line. The leader values at the two step sizes differ by at most 0.0071 at $x \in \{0.5, 1, 2, 5\}$: a drift near the bound makes a boundary layer of width about $\sigma_0^2/(2\bar{\mu}) = 0.0067$ at the exit, which both grids resolve only coarsely.

9. Conclusion

We have given a viscosity-solution treatment of the continuous-time Stackelberg game with a controlled drift and a jump-diffusion follower, in the feedback sense of Definition 1. Existence of an equilibrium was proved in §4 by a Schauder–Tychonoff argument on the coupling map (Theorem 1), under a continuity hypothesis on the best responses. Verification was proved in §5 when the follower's exit region is a half-line (Theorem 2), and that structure was shown for the put payoff under a non-increasing drift (Lemma 2). Section 6 split the jump term off each half-step of the alternating algorithm and proved a sup-norm contraction with rate at most $\lambda/(q + \lambda)$ for the sweeps (Theorem 3 (i)), while the outer iteration reduces to a scalar threshold map whose rate depends on the payoffs (Theorem 3 (ii)). For Kou jumps the follower's best response to a constant drift is in closed form (Proposition 1), resting on the Wiener–Hopf factors and the perpetual-put theorem of [6], and §8 computed both half-steps on a grid: the follower half-step reproduces the closed form at second order with sweeps that contract at rates just below the bound, and the leader half-step obeys the comparison theorem from two starts and against constant-drift sub-solutions.

Monotone stopping was the technical lever throughout: it reduces the follower's free boundary to a single point, it is what the comparison argument of Theorem 2 needs, and it turns the outer iteration into a map on one number.

9.1 Open directions

1. **The outer iteration.** Running the threshold iteration for the Kou example and measuring the map Γ and its rate ρ_Γ ; a bound on ρ_Γ in terms of the payoffs would turn Theorem 3 (ii) into an unconditional uniqueness result.
2. **Commitment.** The bilevel problem of Remark 2, in which the leader optimises through the follower's reaction map; for constant rules Proposition 1 gives that map in closed form.
3. **A continuum of followers.** Replacing the single follower by a continuum couples the leader's HJB equation to a Fokker–Planck equation for the followers who have not exited, in the spirit of Carmona and Delarue [3] and Cardaliaguet [14].
4. **Two-sided exit regions and partial information.** Payoffs that produce two thresholds break Lemma 2 and Theorem 2 (Remark 5), and a leader who observes the follower's position imperfectly needs filtering for jump-diffusions [2].



A. Algorithm pseudocode

Algorithm 1 — Alternating fixed-point iteration with jump-split sweeps

```

1  Input:
2    Kou parameters (sigma_0, lambda, p, eta_+, eta_-), discount  $q > 0$ 
3    Leader payoff (alpha, beta, gamma), drift bound  $\mu_{\text{bar}}$ , salvage kappa
4    Follower payoff  $g(x) = (K - x)_+$ 
5    Uniform grid  $x_1 < \dots < x_N$  on  $[x_{\text{min}}, x_{\text{max}}]$ , step  $h$ 
6    Tolerances  $\eta$  (sweeps) and  $\eta_b$  (threshold); initial threshold  $b_0$ 
7
8  Output:
9    Threshold  $b^*$ , drift rule  $\mu^*(x_i)$ , values  $V_L(x_i)$ ,  $V_F(x_i)$ 
10
11 For  $n = 0, 1, 2, \dots$ :
12   (a) Leader half-step against  $S_n = (-\infty, b_n]$ :
13      $u := \text{kappa} * g$  on  $S_n$ , any start off  $S_n$ .
14     Repeat sweeps:
15        $f := \text{lambda} * K_{\text{hat}}(u)$  (Appendix B;  $u = h_L$  on  $S_n$ )
16       solve  $\sup_{\mu} [D^{\mu} w + \pi_L(x, \mu)] + f - (q + \text{lambda}) w = 0$ 
17         off  $S_n$ ,  $w = h_L$  on  $S_n$ ,
18         by policy iteration over  $\mu = \text{clip}((\text{beta} + w')/\text{gamma})$ 
19       stop when  $\max_i |w(x_i) - u(x_i)| < \eta$ ;  $u := w$ 
20        $V_L := u$ ;  $\mu_n(x_i) := \text{clip}((\text{beta} + V_L'(x_i))/\text{gamma}, \pm \mu_{\text{bar}})$ 
21   (b) Follower half-step under the drift  $\mu_n$ :
22      $u := g$ .
23     Repeat sweeps:
24        $f := \text{lambda} * K_{\text{hat}}(u)$  (Appendix B;  $u = g$  below
25       grid)
26       solve  $\min [ (q + \text{lambda}) w - D^{\mu_n} w - f, w - g ] = 0$ 
27       by policy iteration: on rows marked "stop"  $w = g$ ,
28       on the others the tridiagonal equation; re-mark
29       "stop" where  $w - g \leq (q + \text{lambda}) w - D w - f$ 
30       stop when  $\max_i |w(x_i) - u(x_i)| < \eta$ ;  $u := w$ 
31        $V_F := u$ ;  $b_{n+1} := \text{last grid point where } V_F = g$ 
32   (c) If  $|b_{n+1} - b_n| < \eta_b$ : return  $b_{n+1}$ ,  $\mu_n$ ,  $V_L$ ,  $V_F$ .
33
34 Cost:  $O(N)$  per policy-iteration step and per jump evaluation;
35 about  $\log(1/\eta) / \log(1/\rho)$  sweeps per half-step with
36  $\rho = \text{lambda} / (q + \text{lambda})$  (Theorem 3 (i)); the number of
    outer steps is governed by  $\rho_{\text{Gamma}}$  (Theorem 3 (ii)).

```

B. Jump integrals on the grid

This appendix gives the exact integration of the jump term used in Algorithm 1 and in python/follower_check.py. Let u be piecewise linear on the uniform grid $x_i = x_1 + (i-1)h$. For an exponential jump law of rate η , write $E := e^{-\eta h}$ and



$$c := \frac{1}{h} \left(\frac{1-E}{\eta} - hE \right), \quad a_0 := (1-E) - c, \quad a_1 := c, \quad (\text{B.1})$$

so that $\int_0^h \eta e^{-\eta s} u(x_i \pm s) ds = a_0 u(x_i) + a_1 u(x_{i\pm 1})$, with $a_0, a_1 \geq 0$. The up-jump and down-jump averages $U_i := \int_0^\infty \eta_+ e^{-\eta_+ z} u(x_i + z) dz$ and $D_i := \int_0^\infty \eta_- e^{-\eta_- z} u(x_i - z) dz$ then satisfy the two-term recursions

$$U_i = a_0^+ u_i + a_1^+ u_{i+1} + E_+ U_{i+1}, \quad D_i = a_0^- u_i + a_1^- u_{i-1} + E_- D_{i-1}, \quad (\text{B.2})$$

started from the values of u beyond the grid: $U_N = 0$ when $u = 0$ above the grid, and $D_1 = K - x_1 + 1/\eta_-$ when $u = g = K - x$ below it. Then $\hat{K}u(x_i) = p U_i + (1-p) D_i$. The weights are non-negative and sum to at most one, so the discrete $\hat{K}$ is again a Markov operator and the discrete sweep obeys the bound (6.4) exactly; this is what the largest-ratio column of Table 2 tests.



References

- [1] Bensoussan, A., & Frehse, J. (1984). Nonlinear Elliptic Systems in Stochastic Game Theory. *Journal für die reine und angewandte Mathematik*, 350, 23–67.
- [2] Pham, H. (2009). *Continuous-Time Stochastic Control and Optimization with Financial Applications*. Springer. <https://doi.org/10.1007/978-3-540-89500-8>
- [3] Carmona, R., & Delarue, F. (2018). *Probabilistic Theory of Mean Field Games with Applications*. Springer. <https://doi.org/10.1007/978-3-319-58920-6>
- [4] Mordecki, E. (2002). Optimal Stopping and Perpetual Options for Lévy Processes. *Finance and Stochastics*, 6(4), 473–493. <https://doi.org/10.1007/s007800200070>
- [5] Boyarchenko, S. I., & Levendorskiĭ, S. Z. (2002). *Non-Gaussian Merton-Black-Scholes Theory*. World Scientific. <https://doi.org/10.1142/4955>
- [6] Stochastic Analysis Desk (2023). *Resolvent Operators for One-Dimensional Lévy Processes: A Unified Treatment*. Preprint PR-2026-68980035, <https://zamrik.com/research-items/pr-2026-68980035/>, Zamrik Quant Research.
- [7] Crandall, M. G., Ishii, H., & Lions, P.-L. (1992). User’s Guide to Viscosity Solutions of Second Order Partial Differential Equations. *Bulletin of the American Mathematical Society*, 27(1), 1–67. <https://doi.org/10.1090/s0273-0979-1992-00266-5>
- [8] Protter, P. E. (2004). *Stochastic Integration and Differential Equations*. Springer. <https://doi.org/10.1007/978-3-662-10061-5>
- [9] Schauder, J. (1930). Der Fixpunktsatz in Funktionalräumen. *Studia Mathematica*, 2, 171–180. <https://doi.org/10.4064/sm-2-1-171-180>
- [10] Bensoussan, A., & Lions, J.-L. (1982). *Impulse Control and Quasi-Variational Inequalities*. Gauthier-Villars.
- [11] Howard, R. A. (1960). *Dynamic Programming and Markov Processes*. MIT Press.
- [12] Kou, S. G. (2002). A Jump-Diffusion Model for Option Pricing. *Management Science*, 48(8), 1086–1101. <https://doi.org/10.1287/mnsc.48.8.1086.166>
- [13] Bertoin, J. (1996). *Lévy Processes*. Cambridge University Press.
- [14] Cardaliaguet, P. (2010). Notes on Mean Field Games. *Lecture Notes, Collège de France*.



Zamrik Quant Research

The research arm of Zamrik Quant Lab

About the Desk

Optimal Control & Decision Desk

Zamrik Quant Research


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


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

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


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

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





Fokker–Planck density of the Ornstein–Uhlenbeck process: $\partial_t p = -\partial_x [\kappa(\theta - x)p] + \frac{1}{2}\sigma^2 \partial_{xx} p$



ZAMRIK QUANT RESEARCH

The research arm of Zamrik Quant Lab



PR-2026-46511496
zamrik.com