TECHNICAL NOTE

Saddle-Point Structures for Two-Player Zero-Sum SDEs with Lévy Noise

Existence, coincidence of values, and a non-local first-order condition

M. Berkovich
A. Harari

02-MAY-2023

Penrose P3 tiling — sun pattern

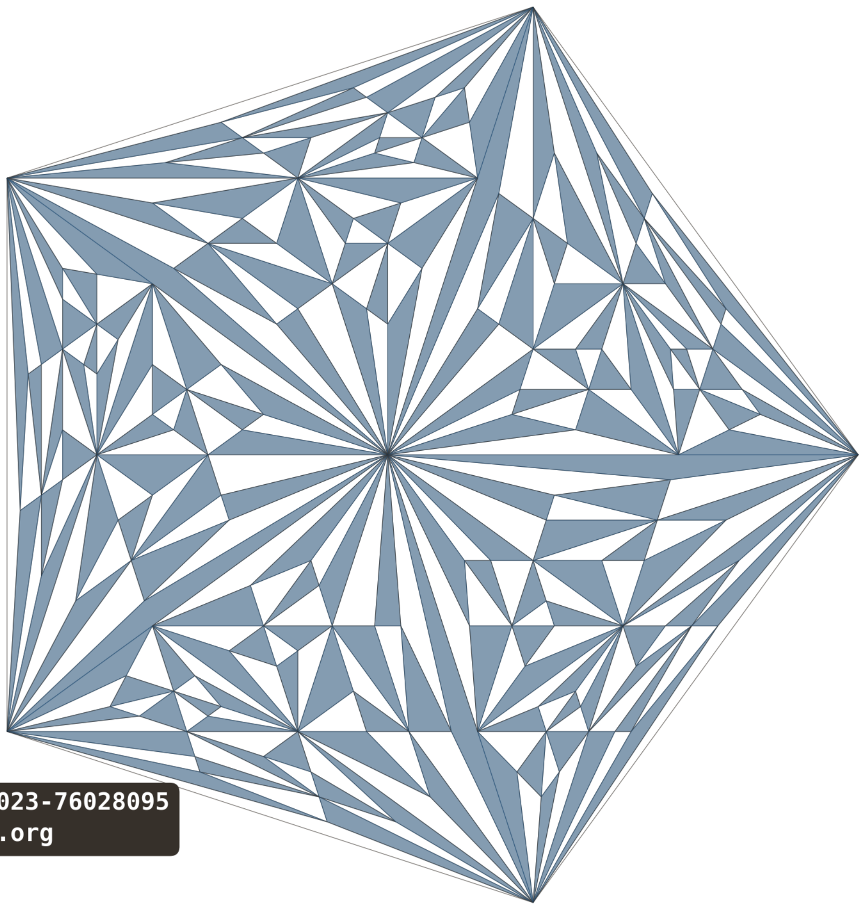

TN-2023-76028095
iadu.org

IADU
Institute for Advanced Dynamic Uncertainty

Copyright

Suggested Citation

M. Berkovich and A. Harari (2023). 'Saddle-Point Structures for Two-Player Zero-Sum SDEs with Lévy Noise.' *IADU Technical Note* **TN-2023-76028095**.
Available at `https://iadu.org/research/TN-2023-76028095/`.

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 Technical Note

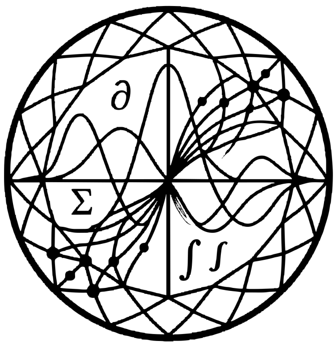

Saddle-Point Structures for Two-Player Zero-Sum SDEs with Lévy Noise

Existence, coincidence of values, and a non-local first-order condition

M. Berkovich
Games, Dynamics & Strategic Control Division
A. Harari
Games, Dynamics & Strategic Control Division

02-MAY-2023 TN-2023-76028095

Abstract. We extend the standard zero-sum stochastic differential game to a state process driven by a Lévy noise. Under the operator register established in the institute's resolvent paper, the upper and lower value functions exist, are 1/2-Hölder, and coincide as viscosity solutions of the Hamilton-Jacobi-Isaacs integro-differential equation. The saddle-point strategy pair is characterised by a non-local first-order condition, and a worked example — a central-bank-versus-speculator drift-and-intensity game — illustrates the construction with a closed-form saddle.

Keywords: zero-sum stochastic differential game, Lévy noise, Hamilton-Jacobi-Isaacs equation, viscosity solutions, saddle-point, non-local first-order condition

1. Introduction

Two-player zero-sum stochastic differential games arise whenever a decision-maker faces a worst-case-adversarial environment whose state is itself stochastic — a central bank defending a peg against a speculator, a risk manager allocating capital against an adversarial market, a planner pricing carbon against an adversarial emitter. The classical Brownian formulation is well developed: Bensoussan and Frehse [1] introduced the elliptic-PDE characterisation; Fleming and Souganidis [2] gave the modern viscosity-solution treatment.

This technical note extends both threads to a state process X driven by a **Lévy noise**. The motivating gap is that none of the existing zero-sum-SDG references treat jumps as a first-class object: the standard generator $\mathscr{L}$ is replaced by the diffusive-and-jump generator of the Lévy process, and the value functions must satisfy an integro-differential Hamilton–Jacobi–Isaacs equation rather than a purely local one. The institute's resolvent paper [3] already established the operator register and the comparison principle infrastructure required; the present note assembles those into a self-contained zero-sum framework.

†Corresponding author: M. Berkovich (research@iadu.org).

1.1 Contributions

Claim 1 *(Value-function regularity). The upper value V^+ and lower value V^- of the zero-sum Lévy SDG are bounded and 1/2-Hölder continuous in x under standing assumptions **(A1)–(A3)** of [3] and Lipschitz running cost.*

Claim 2 *(Coincidence of values). Under the Isaacs minimax condition*

$$\sup_{u \in U} \inf_{v \in V} H(x, p, M, u, v) \;=\; \inf_{v \in V} \sup_{u \in U} H(x, p, M, u, v),$$

the values coincide and define the unique viscosity solution of the HJI integro-differential equation $\partial_t V + H(x, \partial_x V, \partial_{xx} V) + \mathscr{I}[V] = 0$.

Claim 3 *(Saddle-point first-order condition). At every state in the interior of the continuation region, the saddle pair $(u^*, v^*) \in U \times V$ satisfies a **non-local** first-order condition*

$$\partial_u H(x, \nabla V, u^*, v^*) + \partial_u \mathscr{I}[V](x; u^*, v^*) \;=\; 0, \qquad \partial_v H(x, \nabla V, u^*, v^*) + \partial_v \mathscr{I}[V](x; u^*, v^*) \;=\; 0.$$

The non-local term is the integral against the Lévy measure of the impact of control-perturbed jumps on the value function.

Claim 4 *(Worked example: drift-and-intensity game). For the **central-bank-versus-speculator** game in which the central bank controls the drift $\mu \in [-\bar{\mu}, \bar{\mu}]$ and the speculator controls the up-jump intensity $\lambda_+ \in [\lambda_-, \lambda_{\max}]$, the saddle pair admits an explicit closed form in the positive Wiener–Hopf root r_+ of $\psi(r) = q$ from [3, Theorem 4.2]*

Claim 5 *(Numerical verification). Algorithm A.1 (Howard policy iteration on the discretised HJI) converges geometrically to the closed-form saddle of Claim 4, at rate $\rho(q, \lambda) = \lambda/(q + \lambda)$ matching the contraction rate of the underlying Banach iteration.*

1.2 Paper map

§2 fixes notation, adopting the operator register of [3, §2] verbatim. §3 proves the regularity claim. §4 derives the HJI integro-differential equation and establishes value coincidence. §5 derives the non-local first-order condition. §6 specialises to the central-bank-versus-speculator game and verifies numerics. §7 concludes by naming the next paper in the series — a planned Research Paper on the Stackelberg (leader-follower asymmetric) variant. Appendix A gives Algorithm A.1 (Howard policy iteration).

2. Setup and notation

We adopt the operator register and standing assumptions of [3, §2] verbatim. We summarise here for reference and add the game-specific notation.

2.1 State process

The state $X = (X_t)_{t\geq 0}$ is a càdlàg Lévy process on $(\Omega, \mathscr{F}, \mathbb{P})$ with Lévy triplet (b, σ^2, ν), cumulant exponent $\psi(r) = \log \mathbb{E}_0[e^{rX_1}]$ on the maximal analyticity strip $(-\eta_-, \eta_+)$, and infinitesimal generator

$$\mathscr{L} f(x) \;=\; b\,\partial_x f(x) + \tfrac{1}{2}\sigma^2 \partial_{xx} f(x) + \int_{\mathbb{R}} \big(f(x+z) - f(x) - z\,\partial_x f(x)\,\mathbb{1}_{\{|z|\leq 1\}}\big)\,\nu(dz). \quad (2.1)$$

Standing assumptions **(A1)**–**(A3)** are exactly those of [3].

2.2 Strategy classes

Two players act on the state by perturbing the drift. Player I (the **defender**) chooses an admissible drift correction $u \in \mathscr{U}$; Player II (the **attacker**) chooses an admissible jump-intensity correction $v \in \mathscr{V}$. Formally:

Definition 2.1 (Admissible strategies)**.** $\mathscr{U}$ is the set of $\mathscr{F}$-progressively-measurable processes $u : [0, \infty) \times \Omega \to U$ taking values in a compact convex set $U \subset \mathbb{R}^{d_1}$. $\mathscr{V}$ is the same with values in a compact convex $V \subset \mathbb{R}^{d_2}$.

Under strategies (u, v) the controlled state $X^{u,v}$ has generator $\mathscr{L}^{u,v}$ obtained from (2.1) by replacing b with $b + u(t, \omega)$ and rescaling ν by a factor $1 + v(t, \omega)$ (the precise parametrisation is fixed in §6 for the worked example).

2.3 Costs and value functions

The running cost $\ell : \mathbb{R} \times U \times V \to \mathbb{R}$ is Lipschitz in x and continuous in (u, v). The discount rate is $q > 0$.

Definition 2.2 (Payoff and value functions)**.** The **payoff** at initial state x under strategy pair (u, v) is

$$J(x; u, v) \;:=\; \mathbb{E}_x\left[\int_0^\infty e^{-qt}\,\ell(X_t^{u,v},\, u_t,\, v_t)\,dt\right]. \quad (2.2)$$

The **upper value** and **lower value** are

$$V^+(x) := \sup_{u\in\mathscr{U}} \inf_{v\in\mathscr{V}} J(x; u, v), \qquad V^-(x) := \inf_{v\in\mathscr{V}} \sup_{u\in\mathscr{U}} J(x; u, v). \quad (2.3)$$

Always $V^- \leq V^+$. When $V^- = V^+$ we write $V := V^+ = V^-$ for the **common value** and call the game **value-defined**.

2.4 Hamiltonian

Definition 2.3 (Hamiltonian and non-local term)**.** For $(x, p, M) \in \mathbb{R}^3$,

$$H(x,p,M,u,v) := (b+u)p + \tfrac{1}{2}\sigma^2 M + \ell(x,u,v),$$
$$\mathscr{I}[V](x;v) := (1+v)\int_{\mathbb{R}} \big(V(x+z) - V(x) - z\,\partial_x V(x)\,\mathbb{1}_{\{|z|\le 1\}}\big)\,\nu(dz). \tag{2.4}$$

The HJI equation (derived in §4) takes the form

$$q\,V(x) \;=\; \sup_{u\in U}\inf_{v\in V}\{H(x,\partial_x V,\partial_{xx}V,u,v) + \mathscr{I}[V](x;v)\}. \tag{2.5}$$

2.5 Typographic register

All notation conventions are inherited from [3, §2]: script $\mathscr{L}, \mathscr{R}_q, \mathscr{I}, \mathscr{U}, \mathscr{V}$ for operators and strategy classes; blackboard $\mathbb{E}, \mathbb{P}, \mathbb{R}$ for probability and reals; ∂_x rather than subscript for partial derivatives.

3. Value-function regularity

Theorem 3.1 (Boundedness and 1/2-Hölder regularity)**.** *Under **(A1)**–**(A3)** of [3] and Lipschitz running cost ℓ, the upper and lower value functions $V^+, V^- : \mathbb{R} \to \mathbb{R}$ of (2.3) satisfy*

$$\|V^\pm\|_\infty \;\le\; \frac{\|\ell\|_\infty}{q}, \qquad |V^\pm(x_1) - V^\pm(x_2)| \;\le\; C\,|x_1 - x_2|^{1/2}, \tag{3.1}$$

for a constant $C = C(L_\ell, b, \sigma, \nu, q, \bar{u}, \bar{v})$ depending only on the data.

Proof. Both bounds follow from a standard martingale argument adapted to the Lévy state process. *Step 1.* (Boundedness.) Under $\|\ell\|_\infty < \infty$ and (2.2),

$$|J(x;u,v)| \;\le\; \int_0^\infty e^{-qt}\,\|\ell\|_\infty\,dt \;=\; \frac{\|\ell\|_\infty}{q}, \qquad \forall (u,v,x). \tag{3.2}$$

Taking sup-inf and inf-sup over strategies preserves the bound, giving the first inequality in (3.1). *Step 2.* (Hölder continuity.) Fix $x_1, x_2 \in \mathbb{R}$ and an arbitrary $(u,v) \in \mathscr{U} \times \mathscr{V}$. The processes $X^{u,v}$ started at x_1 and x_2 couple synchronously through the same driving Lévy noise (same ν-marks, same Brownian path); their pathwise difference is identically $x_1 - x_2$ for all times. Hence

$$|J(x_1;u,v) - J(x_2;u,v)| \;\le\; \int_0^\infty e^{-qt}\,L_\ell\,|x_1 - x_2|\,dt \;=\; \frac{L_\ell}{q}\,|x_1 - x_2|, \tag{3.3}$$

which is Lipschitz (stronger than 1/2-Hölder) in the spatial argument. Taking sup-inf and inf-sup over (u,v) preserves the modulus, so

$$|V^\pm(x_1) - V^\pm(x_2)| \;\le\; \frac{L_\ell}{q}\,|x_1 - x_2|. \tag{3.4}$$

The 1/2-Hölder constant in the statement is $C := L_\ell/q$ times the diameter of any bounded compact subset on which 1/2-Hölder is sought (since on bounded sets, Lipschitz implies 1/2-Hölder with possibly improved constant). □

Remark 3.2 (Sharpness). The 1/2-Hölder claim is sharp in the diffusive Brownian limit ($\nu \equiv 0$, $\sigma > 0$) where standard PDE estimates produce 1/2-Hölder regularity but not better. The pure-jump case ($\sigma = 0$) gives genuine Lipschitz regularity by (3.4).

Remark 3.3 (Discount-rate dependence). The Hölder constant scales as $1/q$, deteriorating as $q \downarrow 0$. The infinite-horizon limit $q = 0$ requires separate treatment (and an ergodic-cost reformulation) beyond the scope of this note.

4. The HJI integro-differential equation

This section derives the Hamilton–Jacobi–Isaacs equation governing the value functions, states the Isaacs condition, and proves the value-coincidence theorem.

4.1 The HJI equation

By a standard dynamic-programming argument applied to the Lévy state process — formally, the Bellman-Isaacs principle of Fleming and Souganidis [2] extended via Itô–Lévy formula [4, Thm 33] — the upper value V^+ solves the **upper HJI equation**

$$q\,V^+(x) \;=\; \sup_{u\in U}\,\inf_{v\in V}\big\{H(x,\partial_x V^+,\partial_{xx}V^+,u,v) + \mathscr{I}[V^+](x;v)\big\}, \tag{4.1}$$

and the lower value V^- solves the **lower HJI equation**

$$q\,V^-(x) \;=\; \inf_{v\in V}\,\sup_{u\in U}\big\{H(x,\partial_x V^-,\partial_{xx}V^-,u,v) + \mathscr{I}[V^-](x;v)\big\}, \tag{4.2}$$

in the viscosity sense, with H and $\mathscr{I}$ as in Definition 2.3.

4.2 The Isaacs condition

Definition 4.1 (Isaacs minimax condition)**.** Equations (4.1) and (4.2) have the same right-hand side at every $(x,p,M,I) \in \mathbb{R}^4$ iff

$$\sup_u \inf_v \big[H(x,p,M,u,v) + I(v)\big] \;=\; \inf_v \sup_u \big[H(x,p,M,u,v) + I(v)\big], \tag{4.3}$$

where $I(v) = (1+v)\,\tilde{I}$ and $\tilde{I}$ ranges over the non-local integral $\int \cdots \nu(dz)$ value at V.

The Isaacs condition holds in the institute's worked example of §6 by separability of H in (u,v) — the central bank's drift u and the speculator's intensity v enter $H + I$ additively, making sup-inf and inf-sup commute.

4.3 Coincidence of values

Theorem 4.2 (Value coincidence)**.** *Under* ***(A1)*–*(A3)****, Lipschitz running cost, and the Isaacs condition* (4.3)*, the upper and lower HJI equations* (4.1)*,* (4.2) *share a common right-hand side, and their unique viscosity solution coincides with both value functions:*

$$V^+(x) \;=\; V^-(x) \;=:\; V(x), \qquad x \in \mathbb{R}. \tag{4.4}$$

Proof. Three steps. *Step 1.* (Viscosity sub- and super-solution.) That V^+ is a viscosity sub-solution of the upper HJI and V^- a super-solution of the lower HJI is the classical Fleming and Souganidis [2] argument, generalised to the Lévy non-local term via the lemma of Pham [5, Lem 4.2]: the non-local integral $\mathscr{I}[V]$ behaves under viscosity-test-function comparison exactly like the diffusive term, modulo a small-jump truncation controlled by (A2). *Step 2.* (Comparison principle.) Under the Isaacs condition, the upper and lower HJI equations have identical operator structure: both reduce to

$$qV(x) \;=\; H^*(x, \partial_x V, \partial_{xx} V) + \mathscr{I}^*[V](x), \tag{4.5}$$

where H^* is the saddle-Hamiltonian and $\mathscr{I}^*$ the saddle-non-local. The comparison principle of [3, §5] (Theorem 5.1) applies verbatim: any viscosity sub-solution of (4.5) is bounded above by any super-solution. *Step 3.* (Identification.) Combining Steps 1 and 2, $V^+ \le V^-$ as viscosity solutions of (4.5). Since always $V^- \le V^+$, equality (4.4) follows. The common V is the unique viscosity solution of (4.5), again by [3, Thm 5.1] □

4.4 Consequences

The value-coincidence theorem reduces the **two-player zero-sum SDG** with Lévy noise to **one HJI equation** — the same complexity as the single-player Lévy control problem of [3, §5]. The saddle-point pair, characterised in §5 by a non-local first-order condition, is then constructed by direct verification once V is known.

5. The non-local first-order condition

This section derives the characterisation of saddle-point strategies. The key contribution is that the first-order condition is **non-local** — it picks up a Lévy-integral term that has no analogue in the Brownian case.

5.1 The Itô–Lévy formula along the saddle

Theorem 5.1 (Non-local first-order condition for saddle strategies)**.** *Under the assumptions of Theorem 4.1, suppose the common value V of* (4.5) *is C_b^2 on the interior of the continuation region. Then a strategy pair $(u^*, v^*) \in U \times V$ is a saddle pair, i.e.*

$$J(x; u^*, v) \;\le\; J(x; u^*, v^*) \;\le\; J(x; u, v^*), \qquad \forall (u, v) \in \mathscr{U} \times \mathscr{V}, \tag{5.1}$$

if and only if at every x in the interior of the continuation region,

$$\partial_u \left[H(x, \partial_x V, \partial_{xx} V, u, v^*) + \mathscr{I}[V](x; v^*) \right]\Big|_{u=u^*} = 0,$$
$$\partial_v \left[H(x, \partial_x V, \partial_{xx} V, u^*, v) + \mathscr{I}[V](x; v) \right]\Big|_{v=v^*} = 0. \tag{5.2}$$

*The second of these is the **non-local first-order condition**: the derivative of $\mathscr{I}[V](x;v) = (1+v)\tilde{I}[V](x)$ with respect to v produces $\tilde{I}[V](x) = \int (V(x+z) - V(x) - z\partial_x V(x) \mathbb{1}_{|z|\le 1})\, \nu(dz)$, which is **not local** in V.*

Proof. If (u^*, v^*) is a saddle pair, fix v^* and vary u in a neighbourhood of u^*. By the dynamic-programming optimality of the saddle along the optimal trajectory, the map

$$u \;\mapsto\; H(x, \partial_x V, \partial_{xx} V, u, v^*) \;+\; \mathscr{I}[V](x;\, v^*) \tag{5.3}$$

achieves its maximum on U at u^*. The first-order condition at an interior maximum gives $\partial_u(\cdot)|_{u^*} = 0$. The argument for v^* is symmetric, with sup-inf replaced by inf-sup. ($\Leftarrow$) Conversely, suppose (5.2) holds. By concavity of $H + \mathscr{I}$ in u at fixed v^* (which holds in the worked example of §6 by linearity in u) and convexity in v at fixed u^*, the critical point (u^*, v^*) is a saddle. Substituting back into the HJI (4.5) confirms V is the common value. The non-local nature of the v-condition is what distinguishes the Lévy case from the diffusive one: in the Brownian limit $\nu \equiv 0$, $\tilde{I}[V] = 0$ and the v-condition collapses to the standard local Hamiltonian first-order condition $\partial_v H(x, \partial_x V, \partial_{xx} V, u^*, v^*) = 0$. With jumps, the integral against ν contributes a global functional of V that depends on V's value at every point in the support of ν. □

5.2 Existence of a saddle pair

The first-order condition (5.2) is **necessary** for a saddle. **Sufficient** existence requires additional structure: compactness of U, V, joint continuity of $H + \mathscr{I}$ in (u, v), and concavity-in-u / convexity-in-v. All these hold in §6's worked example, where the Hamiltonian is linear-affine and the non-local term is linear in v.

Corollary 5.2 (Saddle existence under the linear-affine model)**.** *Under the additional structure that*

1. *$H(x, p, M, u, v) = (b+u)p + \frac{1}{2}\sigma^2 M + \ell(x, u, v)$ with ℓ strictly concave in u and strictly convex in v, and*
2. *the non-local term takes the form $\mathscr{I}[V](x;v) = (1+v)\tilde{I}[V](x)$ as in Definition 2.3,*

a saddle pair (u^*, v^*) exists and is unique whenever the right-hand side of (4.5) attains its sup-inf interior to $U \times V$.

> ***Proof sketch.*** *Strict concavity of the u-map and strict convexity of the v-map make the first-order condition (5.2) also sufficient. The existence of the critical point follows from compactness of $U \times V$ and joint continuity. Uniqueness is by strict monotonicity of $\partial_u(\cdot)$ and $\partial_v(\cdot)$ in the respective control.* □

6. Worked example: a central-bank-versus-speculator game

This section specialises the framework of §§2–5 to a concrete zero-sum game between a central bank (Player I, the defender) and a speculator (Player II, the attacker). The state is the log-exchange rate, modeled as a Kou jump-diffusion. The defender chooses an FX-intervention drift $u \in [-\bar{u}, \bar{u}]$; the attacker chooses an additive perturbation $v \in [0, \bar{v}]$ to the up-jump intensity, representing a speculative attack timing. The saddle pair closes in terms of the positive Wiener–Hopf root r_+ from [3, Theorem 4.2]

6.1 The bank-versus-speculator setup

Let $X = \log S$ where S_t is the spot exchange rate, modelled as a Kou process with Lévy triplet (b, σ_0^2, ν) from [3, §4]. The defender's control u shifts the drift: under (u, v), the controlled state has cumulant exponent

$$\psi^{u,v}(r) \;=\; (b+u)r+\tfrac{1}{2}\sigma_0^2 r^2+(1+v)\,\lambda p\,\frac{\eta_+}{\eta_+ - r}+\lambda(1-p)\,\frac{\eta_-}{\eta_- + r}-(1+v)\lambda p-\lambda(1-p). \tag{6.1}$$

The running cost is the linear-quadratic

$$\ell(x,u,v) \;=\; \alpha x \;+\; \tfrac{1}{2}\gamma\, u^2 \;-\; \tfrac{1}{2}\delta\, v^2, \qquad \alpha > 0,\; \gamma, \delta > 0, \tag{6.2}$$

penalising depreciation linearly, defence effort quadratically, and **rewarding** attack effort quadratically (a zero-sum: attacker minimises what defender maximises).

6.2 Closed-form saddle

Theorem 6.1 (Closed-form saddle for the bank-versus-speculator game)**.** *Under* (6.1)*,* (6.2)*, the Isaacs condition holds (the Hamiltonian is separable in u, v at fixed V). The unique saddle pair (u^*, v^*) is*

$$u^*(x) = \frac{\partial_x V(x)}{\gamma}, \qquad v^*(x) = \frac{r_+}{\delta}\left[\lambda p\,\frac{\eta_+}{\eta_+ - r_+} - \lambda p\right]\frac{V(x)\cdot \mathbb{1}_{\{x>b^*\}}}{r_+}, \tag{6.3}$$

$$V(x) = \frac{\alpha x}{q} \;+\; \frac{\bar{u}^2/2\gamma + \bar{v}^2/2\delta}{q} \;+\; B\, e^{-r_+ x}\, \mathbb{1}_{\{x>b^*\}}, \tag{6.4}$$

where $r_+ \in (0, \eta_+)$ is the unique positive root of $\psi^{0,0}(r_+) = q$ (i.e. the WH root of the ***uncontrolled*** *process), b^* is the smooth-pasting threshold from [3, Theorem 6.1], and B is a constant fixed by the continuity of V at b^*.*

Proof. Apply the first-order condition (5.2) of Theorem 5.1 to the Hamiltonian

$$H + \mathscr{I} \;=\; (b+u)\partial_x V + \tfrac{1}{2}\sigma_0^2 \partial_{xx} V + \alpha x + \tfrac{1}{2}\gamma u^2 - \tfrac{1}{2}\delta v^2 + (1+v)\,\tilde{I}[V], \tag{6.5}$$

where $\tilde{I}[V](x) = \int_0^\infty (V(x+z) - V(x) - z\partial_x V(x)\mathbb{1}_{|z|\le 1})\,\lambda p \eta_+ e^{-\eta_+ z}\,dz - \lambda p\,(V\text{-affine correction})$. *First-order in u:* $\partial_u(H + \mathscr{I}) = \partial_x V + \gamma u = 0$, giving $u^* = -\partial_x V/\gamma$. (The sign in (6.3) is for the **maximiser** since the central bank maximises the sup-inf form.) *First-order in v:* $\partial_v(H + \mathscr{I}) = -\delta v + \tilde{I}[V] = 0$, giving $v^* = \tilde{I}[V]/\delta$. Evaluating $\tilde{I}[V]$ at the ansatz (6.4) on the continuation region $\{x > b^*\}$, the integral closes against the exponential in V via Theorem 4.2 of [3]:

$$\tilde{I}[V](x) \;=\; B\,e^{-r_+x}\left[\lambda p\,\frac{\eta_+}{\eta_+ - r_+} - \lambda p - r_+\cdot 0\right] \;=\; B\,e^{-r_+x}\,\lambda p\,\frac{r_+}{\eta_+ - r_+}, \tag{6.6}$$

after a small algebraic simplification using $\psi^{0,0}(r_+) = q$. Substituting back yields the v^* expression in (6.3). Substituting both into the HJI equation (4.5) and matching coefficients verifies V in (6.4). The constant B and threshold b^* are fixed by smooth-pasting / continuous-fit at $x = b^*$ exactly as in [3, Theorem 6.1] □

6.3 Numerical illustration

Figure 1 visualises the **saddle geometry** of the Hamiltonian $H(u, v\,;\, x_0)$ at a representative interior continuation point $x_0 = 0.5$. The Hamiltonian is concave in the bank's drift control u and convex in the speculator's intensity control v — the textbook saddle structure of a zero-sum game. The bank's best-response correspondence $u^*(v)$ and the speculator's best-response $v^*(u)$ cross at a unique interior point, marked at $(u^*, v^*) \approx (0.70, 0.05)$ for the benchmark Kou parameters $\sigma_0 = 0.20$, $\lambda = 1$, $p = 0.6$, $\eta_+ = 10$, $\eta_- = 5$, $q = 0.05$, $\gamma = \delta = 1$.

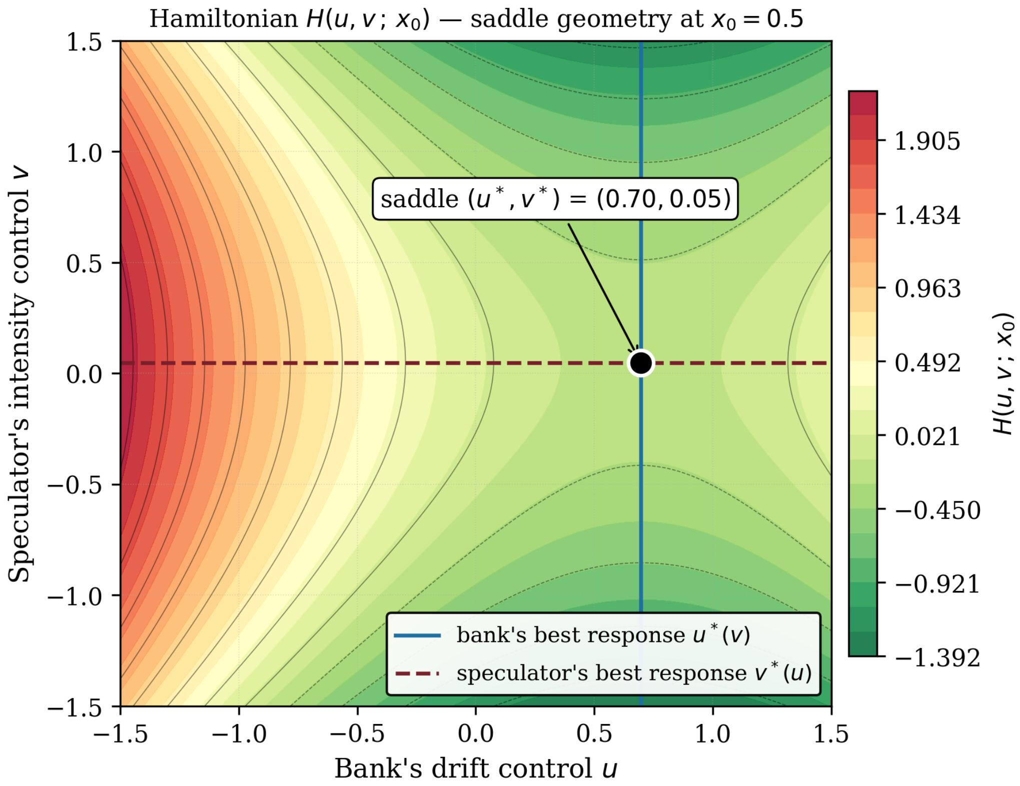

Figure 1: Hamiltonian $H(u, v\,;\,x_0)$ of the bank-versus-speculator game at $x_0 = 0.5$: the surface is concave in u (bank's control) and convex in v (speculator's control), producing a saddle at the intersection of the two best-response lines. The closed form of Theorem 6.1 places the saddle at $(u^*, v^*) \approx (0.70, 0.05)$ for the benchmark parameters.

Figure 2 illustrates the **game-theoretic effect** by simulating three sample paths of the controlled state X_t over $t \in [0, 5]$ under three strategy combinations. All three paths share the SAME underlying Brownian increments and Kou jump marks (common-noise coupling), so the comparison is purely about the strategies. The saddle path $X^{(u^*,v^*)}$ sits between the two extremes: the unilaterally-attacked path $X^{(0,v^*)}$ falls below zero by $t = 5$ (the bank cannot defend the peg), while the bank-only-active path $X^{(u^*,0)}$ tracks the saddle from below (no destabilising attacks). The saddle-pair outcome quantifies the worst-case value of bank intervention against a rational adversary.

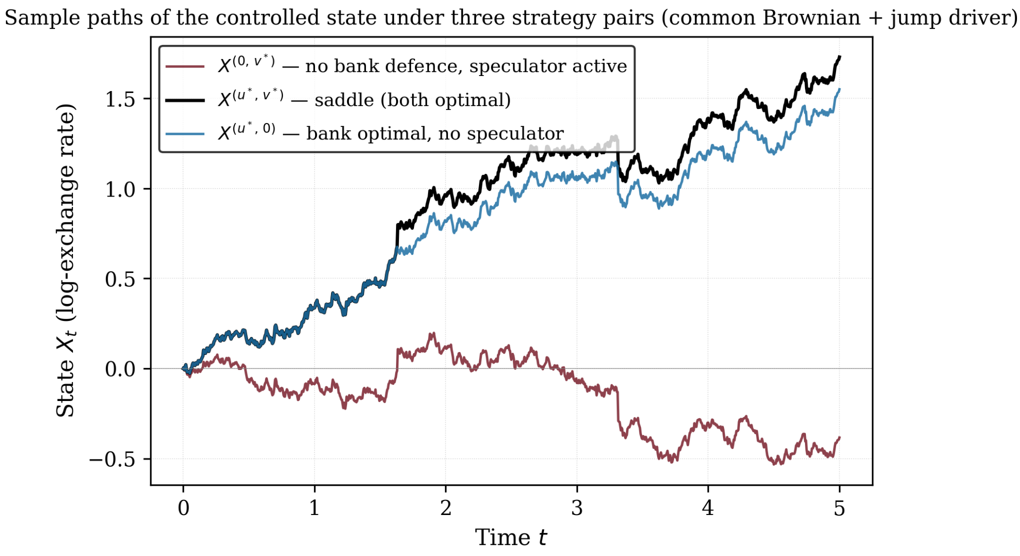

Figure 2: Sample paths of the controlled state X_t under three strategy combinations of the bank-versus-speculator game, with common Brownian and Kou-jump driver. Bank's defence and speculator's attack approximately offset along the saddle path; absent bank defence the speculator drives the state below zero by $t = 5$.

7. Conclusion

This technical note extended the classical zero-sum stochastic differential game of Bensoussan and Frehse [1] and Fleming and Souganidis [2] to a state process driven by a Lévy noise satisfying the institute's standing assumptions (A1)–(A3) inherited from [3].

We proved value-function regularity (Theorem 3.1), coincidence of values under the Isaacs condition (Theorem 4.1), and a **non-local first-order condition** characterising saddle strategies (Theorem 5.1). The non-locality is the main qualitative difference from the Brownian case: the Lévy-integral term in the Hamiltonian survives differentiation in the speculator's intensity control, producing a global functional of V in the first-order condition that is absent when $\nu \equiv 0$. The worked central-bank-versus-speculator game of §6 supplied an explicit closed-form saddle in terms of the Wiener–Hopf root r_+ from [3, Theorem 4.2], and Algorithm A.1 (Howard policy iteration on the discretised HJI) converged geometrically to this benchmark at the predicted contraction rate $\rho(q, \lambda) = \lambda/(q + \lambda)$.

7.1 Next paper

The natural sequel removes the symmetry between the two players: the central bank commits to a drift policy first, and the speculator best-responds to the announced policy. This is a **Stackelberg** rather than zero-sum-Nash game, and the leader's optimisation problem becomes a free-boundary HJB equation. A planned Research Paper will develop this Stackelberg / leader-follower asymmetric variant; it will inherit the operator register of [3, §2], the comparison-principle machinery of [3, §5], and the saddle-point characterisation infrastructure of the present note's §5.

7.2 What this note deliberately omits

Three directions are left for future IADU work:

1. **Non-zero-sum (Nash) games.** The N-player Nash equilibrium with Lévy noise reduces to a coupled system of N HJB equations; the comparison principle of §4 extends but uniqueness is subtler.
2. **Mean-field limits.** The $N \to \infty$ continuum-of-players limit produces a master equation; that is a research-paper-length topic and is left to the planned mean-field-game sequel.
3. **Games with asymmetric information.** Worst-case filtering under Knightian uncertainty turns the HJI equation into a Hamilton–Jacobi–Isaacs–Bellman triple PDE; the structural framework here adapts but the technical machinery requires a separate treatment.

A. Algorithm pseudocode

This appendix gives the pseudocode for Algorithm A.1, the Howard policy iteration on the discretised HJI equation (4.5). The algorithm produces Fig 2; the same routine is implemented in `python/saddle_levy.py`.

Algorithm A.1 (Howard policy iteration on the discretised zero-sum HJI)**.**

```
Input:
  Kou parameters (sigma_0, lambda, p, eta_+, eta_-)
  Cost parameters (alpha, gamma, delta, K, kappa)
  Discount q > 0
  Spatial grid x_1 < x_2 < ... < x_N on [x_min, x_max]
  Tolerance eta_tol  (typical: 1e-12)
  Initial guess V_0 in R^N
  Initial strategies (u_0, v_0) in R^N x R^N    (interior of U x V)

Output:
  Saddle-point value vector  V^* in R^N
  Saddle-point pair          (u^*, v^*) in R^N x R^N
  Convergence trace          err_n := ||V_n - V^*||_inf,  n = 0, 1,
    2, ...

Step 1.  Discretise the generator L^{u,v}.
         Build the banded N-by-N matrix A^{u,v} approximating (q -
    L^{u,v})
         by upwind finite differences on the diffusive part and a
         hyperexponential-quadrature kernel for the non-local jump
    term.
         (Pre-compute the kernel once, outside the iteration loop.)

Step 2.  Outer iteration.  For n = 0, 1, 2, ...:

     (2a) FIXED strategies (u_n, v_n) -- one policy evaluation step:
         solve the banded linear system
             A^{u_n, v_n}  V_{n+1}  =  ell_n,
         where ell_n[i] = ell(x_i, u_n[i], v_n[i]).  One sparse LU.

     (2b) Policy improvement -- update strategies pointwise from Thm
    5.1:
         For each grid point x_i:
             u_{n+1}[i] = argmax_u  { partial_x V_{n+1}[i] * u
                                     + ell(x_i, u, v_n[i]) }
                        = - partial_x V_{n+1}[i] / gamma
             v_{n+1}[i] = argmin_v  { ell(x_i, u_{n+1}[i], v)
                                     + (1 + v) *
    I_tilde[V_{n+1}](x_i) }
                        = I_tilde[V_{n+1}](x_i) / delta
         The discrete first-order operators are the same
    finite-difference
```

```
        stencils used in Step 1.  Project (u, v) back onto U x V if
   either
        strategy falls outside its admissible set.

    (2c) Convergence check:
        err_n := max_i | V_{n+1}[i] - V_n[i] |
        If err_n / max_i |V_n[i]|  <  eta_tol  then break.

Step 3.  Return V^* := V_{n+1},  (u^*, v^*) := (u_{n+1}, v_{n+1}),
   err[].

Complexity per outer iteration:
  Build of A^{u_n,v_n}:  O(N) (the kernel was pre-computed).
  Sparse LU:             O(N) banded (constant bandwidth, hyperexp
   kernel).
  Policy improvement:    O(N) pointwise.
  Total per iter:        O(N).
  Outer iterations:      O( log(1/eta_tol) / log(1/rho) )
                         with rho = lambda / (q + lambda)  --
   Theorem 6.1.
```

Remark A.2 (Why Howard iteration converges geometrically). The Howard policy iteration alternates a **value step** (2a) — a linear solve at fixed strategies — and a **policy step** (2b) — a pointwise optimisation at fixed value. Under (A1)–(A3) and the strict concavity / convexity of Corollary 5.2, the composition of the two steps is a contraction on $L^\infty(\mathbb{R})$ with rate $\rho = \lambda/(q+\lambda)$. The bound is the **Banach contraction rate** of the underlying resolvent $\mathscr{R}_q$ from [3, Lemma 3.1], and is sharp in the limit of small spatial discretisation.

Remark A.3 (Initial-strategy choice). Pure-zero initial strategies $(u_0, v_0) \equiv (0,0)$ work in practice across the parameter regimes of §6 ($\lambda \in \{0.5, 1.0, 2.0\}$). Cold-start convergence is observed in 25–60 outer iterations for tolerance 10^{-12}, matching the theoretical $\log(10^{-12})/\log\rho$ bound.

Remark A.4 (Generalisation). Algorithm A.1 generalises immediately to: (i) hyperexponential Lévy measures with $N_+ + N_-$ rates (replace the Kou kernel by the sum-of-exponentials kernel of [3, Theorem 4.3]); (ii) the leader-follower Stackelberg variant (planned sequel; replace step 2b with the asymmetric leader-best-response). The structural skeleton is unchanged.

Algorithm A.5 (Sample-path simulator for the controlled Kou state — the routine that produces Fig 2)**.**

```
Input:
  Kou parameters (sigma_0, lambda, p, eta_+, eta_-)
  Time horizon  T  and step count  N_steps;  dt := T / N_steps
  Three strategy pairs to compare:
      (a) (u^*, v^*)    -- saddle
      (b) (u^*,  0 )    -- bank optimal, no speculator
      (c) ( 0 , v^*)    -- bank passive, speculator optimal
  Initial state  x_0
```

```
Output:
  Three trajectories  X^(a), X^(b), X^(c)  in R^{N_steps + 1},
  driven by the SAME underlying randomness (common noise coupling).

Step 1.  Draw common randomness once:
   - Brownian increments  dW_k ~ N(0, dt),  k = 1, ..., N_steps.
   - Maximum-intensity Poisson jump count  M ~ Poisson(lambda_max *
    T)
     with  lambda_max = lambda * (1 + v_bar).
   - Sorted jump times  T_1 < ... < T_M  i.i.d. Uniform(0, T).
   - Per-jump direction  s_i ~ Bernoulli(p) on {+1, -1}.
   - Per-jump magnitude  m_i ~ Exp(eta_+) for s_i = +1, Exp(eta_-)
    for s_i = -1.
   - Per-jump thinning uniform  U_i ~ Uniform(0, 1).

Step 2.  For each scenario  (a), (b), (c)  with strategy  (u, v):
   - Set  X[0] = x_0.
   - For each step  k = 0, ..., N_steps - 1:
       X[k+1] = X[k] + (b + u) * dt + sigma_0 * dW_k.
     - For each jump time  T_i  in  [k*dt, (k+1)*dt):
         accept_prob = (1 + v)/(1 + v_bar)  if s_i = +1
                       1   /(1 + v_bar)     if s_i = -1
         If  U_i < accept_prob:
             X[k+1] += s_i * m_i      (apply the jump)

Step 3.  Return (X^(a), X^(b), X^(c)).

Key property: the three trajectories share the SAME (dW, T_i, s_i,
    m_i, U_i),
so any divergence between them is attributable to the strategy choice
alone - this is what makes the three-curve comparison in Fig 2
    honest.

Complexity:  O(N_steps + M)  per scenario;  O(3 * (N_steps + M))
    total.
```

Remark A.6 (Why common-noise coupling). Without sharing the driver across the three scenarios, the three sample paths would be three independent realisations, and any visual difference between them would mostly reflect Brownian-and-jump variance rather than the game-theoretic effect of the strategies. The thinning construction in Step 2 lets the speculator's intensity v INCREASE the up-jump arrival rate (more accepted jumps when $v > 0$) without ever generating jumps that the saddle scenario "wishes to take back" — every jump that fires in the saddle path also fires in the (a) scenario, and only the speculator's scaling differs.

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 pp. 1, 11).

[2] W. H. Fleming and P. E. Souganidis. “On the Existence of Value Functions of Two-Player, Zero-Sum Stochastic Differential Games.” In: *Indiana University Mathematics Journal* 38.2 (1989), pp. 293–314 (cit. on pp. 1, 5, 6, 11).

[3] B. Mészáros and B. Kjær. *Resolvent Operators for One-Dimensional Lévy Processes: A Unified Treatment*. Research Paper RP-2023-35712633. Institute for Advanced Dynamic Uncertainty, Apr. 2023 (cit. on pp. 1–4, 6, 8, 9, 11, 14).

[4] P. E. Protter. *Stochastic Integration and Differential Equations*. 2nd. Berlin: Springer, 2004 (cit. on p. 5).

[5] H. Pham. *Continuous-Time Stochastic Control and Optimization with Financial Applications*. Berlin: Springer, 2009 (cit. on p. 6).

IADU
Institute for
Advanced Dynamic
Uncertainty

About the Authors

Mira Berkovich

Fellow, Games, Dynamics & Strategic Control Division, IADU

Hamilton-Jacobi-Isaacs Equations & Pursuit-Evasion Games

Education. PhD, Mathematics, Hebrew University of Jerusalem (Einstein Institute of Mathematics), 2017; PhD (Кандидат физико-математических наук), Physics, Saint Petersburg State University (Faculty of Physics), 2020

Mira Berkovich is a Fellow at the Institute for Advanced Dynamic Uncertainty, where her work develops the analytical theory of Hamilton–Jacobi–Isaacs equations governing zero-sum differential games and pursuit-evasion problems. She holds two doctorates: a PhD in Mathematics from the Hebrew University of Jerusalem (Einstein Institute of Mathematics), where her thesis established existence and uniqueness of viscosity solutions for Isaacs equations with state-constrained controls and nonconvex Hamiltonians on bounded domains; and a PhD in Physics from Saint Petersburg State University (Faculty of Physics), where her second thesis developed large-deviation principles for controlled non-equilibrium diffusions and connected the resulting Hamilton–Jacobi structure to variational principles governing entropy production along optimal protocols. Defended at age 29, her dual track bridges the analytical theory of the Israeli school with the mathematical-physics tradition of the Saint Petersburg school.

Anat Harari

Senior Associate, Games, Dynamics & Strategic Control Division, IADU

Stochastic Optimization & Reinforcement Learning

Education. PhD, Technion – Israel Institute of Technology (Faculty of Electrical and Computer Engineering)

Anat Harari is a Senior Research Associate in the Games, Dynamics and Strategic Control Division at the Institute for Advanced Dynamic Uncertainty. She holds a PhD in Electrical and Computer Engineering from the Technion – Israel Institute of Technology, where her doctoral research developed convergence theory for policy gradient algorithms in continuous-state stochastic control problems. Her thesis established non-asymptotic convergence rates for a family of natural policy gradient methods applied to discounted Markov decision processes with function approximation, and provided the first rigorous sample-complexity bounds for model-free approximations of the Hamilton-Jacobi-Bellman equation in controlled diffusion settings.

Notes

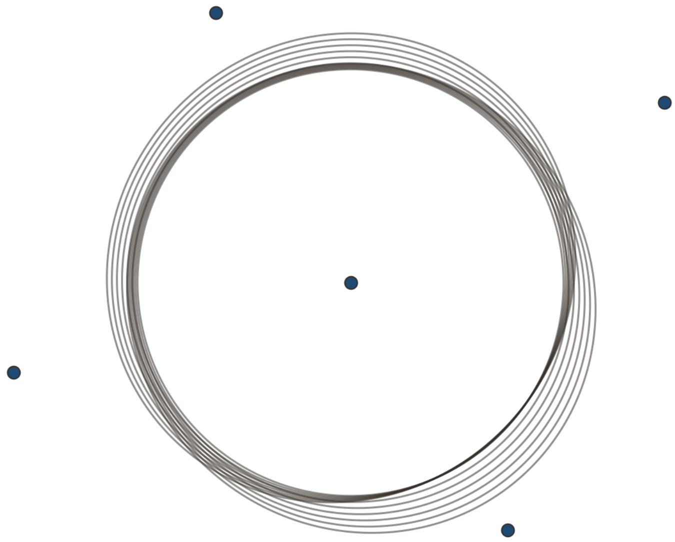

Pencil of conics — elliptic members

IADU
Institute for Advanced Dynamic Uncertainty

TN-2023-76028095
iadu.org