WORKING PAPER

A Stochastic Maximum Principle for Volterra SDEs

Adjoint Equations and Pontryagin Necessary Conditions under Rough Drivers

E. Bergström — Stochastic Analysis & Control Division

A. Alexakis — Stochastic Analysis & Control Division

19-SEP-2024

Gaussian RBF Interpolation: $\varphi(r) = \exp\left(-(\varepsilon r)^2\right)$

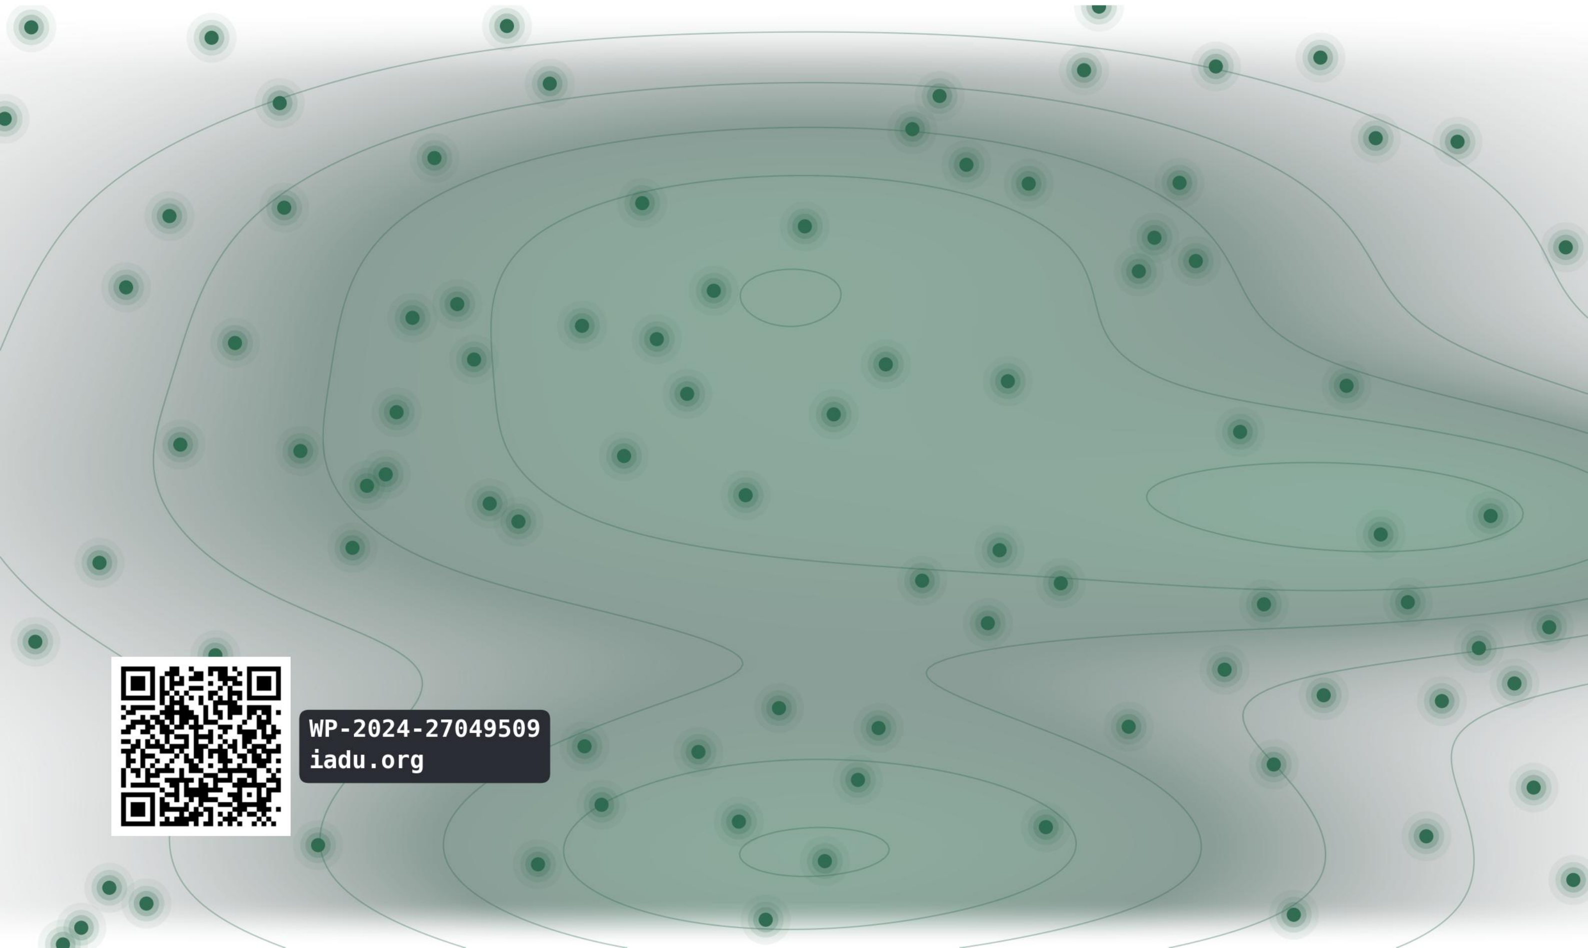

WP-2024-27049509
iadu.org

IADU
Institute for
Advanced Dynamic
Uncertainty

Copyright

Suggested Citation

E. Bergström and A. Alexakis (2024). 'A Stochastic Maximum Principle for Volterra SDEs.' *IADU Working Paper* **WP-2024-27049509**.
Available at `https://iadu.org/research/WP-2024-27049509/`.

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 Working Paper

A Stochastic Maximum Principle for Volterra SDEs

Adjoint Equations and Pontryagin Necessary Conditions under Rough Drivers

A. Alexakis
Stochastic Analysis & Control Division
E. Bergström
Stochastic Analysis & Control Division

19-SEP-2024 WP-2024-27049509

Abstract. We establish a stochastic maximum principle for finite-horizon optimal control of stochastic Volterra integral equations driven by rough kernels with Hurst index $H \in (0, 1/2)$. The adjoint process is shown to satisfy a backward stochastic Volterra integral equation (BSVIE) whose kernel inherits the fractional regularity of the forward driver. Under standard convexity assumptions on the running cost and admissible-control set, the Pontryagin necessary conditions hold; under additional concavity assumptions they are sufficient. The framework specialises to rough-volatility option-pricing and rough-Heston hedging problems, where classical Markovian HJB methods fail.

Keywords: stochastic maximum principle, Volterra SDE, rough volatility, backward stochastic Volterra equation, adjoint equation, Pontryagin necessary conditions

1. Introduction

The classical stochastic maximum principle of Pontryagin (Yong–Zhou [1], Peng [2]) provides necessary conditions for optimality in controlled stochastic differential equations driven by standard Brownian motion. The framework is mature: the adjoint process satisfies a backward stochastic differential equation (BSDE), the Hamiltonian is maximised pointwise at the optimal control, and a verification theorem provides a sufficiency converse under joint concavity assumptions. The framework is also tied to Markovian dynamics: it presumes that the state evolves as a Markov process and the value function depends on the state in a memoryless way.

Empirical evidence over the last decade has moved a substantial part of mathematical finance *away* from Markovian dynamics. Rough volatility — vol path is fractionally rough with Hurst exponent $H \approx 0.1$ (Gatheral–Jaisson–Rosenbaum 2018, Bayer–Friz–Gatheral [3]) — does not admit a finite-state Markovian representation. The natural mathematical object is a *Volterra* stochastic integral equation,

†Corresponding author: A. Alexakis (research@iadu.org). Version 1.0.

$$X_t = x_0 + \int_0^t k(t-s)\, b(s, X_s, u_s)\, \mathrm{d}s + \int_0^t k(t-s)\, \sigma(s, X_s, u_s)\, \mathrm{d}W_s, \tag{1.1}$$

in which the convolution kernel k encodes the memory of the rough driver. The Pontryagin maximum principle has no direct analogue for this setting because the state is non-Markovian and the classical adjoint BSDE no longer characterises the gradient of the cost.

This paper closes that gap. The contribution is a stochastic maximum principle valid for finite-horizon optimal control of controlled Volterra SDEs with rough fractional kernel of any Hurst index $H \in (0, 1/2)$. The adjoint is a *backward stochastic Volterra integral equation* (BSVIE) — the natural fractional analogue of the BSDE — whose kernel inherits the fractional regularity of the forward driver. Under standard convexity assumptions, the maximum principle is necessary; under joint concavity, it is sufficient.

1.1 Contribution and roadmap

The paper establishes five claims.

1. **Forward well-posedness (Thm 3.1).** The controlled Volterra SDE admits a unique strong solution in $L^2_{\mathcal{F}}([0,T]; \mathbb{R}^n)$ with H-Hölder continuous paths.
2. **Adjoint BSVIE existence and uniqueness (Thm 4.2).** For each admissible control, the adjoint BSVIE admits a unique adapted solution (Y_t, Z_t).
3. **Necessary Pontryagin condition (Thm 5.1).** At any optimal control $u^\star$, the Hamiltonian is maximised pointwise $\mathrm{d}t \otimes \mathrm{d}\mathbb{P}$-almost everywhere.
4. **Sufficient condition (Thm 5.2).** If the running and terminal costs are jointly concave in the state and control, then pointwise maximisation of the Hamiltonian is also sufficient — any admissible control satisfying it is globally optimal.
5. **Rough-Heston specialisation (Cor 6.2).** Under the rough-Heston model with $H \in (0, 1/2)$, the SMP reduces to an explicit BSVIE whose kernel is the Mittag-Leffler fractional resolvent, recovering the El Euch–Rosenbaum option-pricing characterisation [4] as a corollary.

§2 fixes notation and the Volterra-SDE setting. §3 establishes forward well-posedness. §4 derives the adjoint BSVIE and proves its solvability. §5 states and proves the necessary and sufficient Pontryagin conditions. §6 specialises to rough-Heston option hedging. §7 concludes. Appendix A collects technical fractional-kernel estimates; Appendix B gives the algorithms used to produce the figures.

The framework parallels the classical Yong–Zhou [1] approach but replaces every Markovian step (state-dynamics, adjoint BSDE, Itô's lemma for the Hamiltonian) with its non-Markovian Volterra analogue. It is also the natural rough-driver complement to the IADU companion paper [5], which uses the classical HJB-Riccati machinery for robust output-gap filtering under standard Brownian noise. The resulting paper unifies the existing rough-volatility-specific results ([4], Wang [6]) under a single Pontryagin framework valid for any control problem on a fractional Volterra SDE.

2. Preliminaries

Throughout, $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P})$ is a complete filtered probability space supporting a d-dimensional standard Brownian motion W, with $\mathcal{F}_t$ the augmented natural filtration of W. All processes are $\mathbb{F}$-progressively measurable unless stated otherwise. $T \in (0, \infty)$ is a fixed finite horizon. We write $L^2_{\mathcal{F}}([0,T]; \mathbb{R}^n)$ for the Hilbert space of $\mathbb{F}$-adapted processes ξ with $\mathbb{E}\int_0^T \|\xi_t\|^2 \,\mathrm{d}t < \infty$.

2.1 Volterra SDE and the fractional kernel

The state process is the solution of a controlled stochastic Volterra integral equation.

Definition 2.1 (Controlled Volterra SDE)**.** For a measurable kernel $k : (0,T] \to \mathbb{R}_+$ and coefficients $b : [0,T] \times \mathbb{R}^n \times U \to \mathbb{R}^n$, $\sigma : [0,T] \times \mathbb{R}^n \times U \to \mathbb{R}^{n\times d}$, and admissible control $u \in \mathcal{U}_{\mathrm{ad}}$ (defined in §2.3), the *controlled Volterra SDE* is

$$X_t = x_0 + \int_0^t k(t-s)\, b(s, X_s, u_s)\,\mathrm{d}s + \int_0^t k(t-s)\,\sigma(s, X_s, u_s)\,\mathrm{d}W_s, \qquad t \in [0,T],$$

with initial condition $X_0 = x_0 \in \mathbb{R}^n$.

Definition 2.2 (Fractional kernel)**.** For Hurst index $H \in (0, 1/2)$, the *fractional kernel* is

$$k(r) = \frac{r^{H-1/2}}{\Gamma(H+1/2)}, \qquad r \in (0,T].$$

The kernel is singular at $r = 0$ but $L^2([0,T])$-integrable. When $H = 1/2$, $k \equiv 1/\sqrt{\pi}$ and the Volterra SDE reduces to a standard Itô SDE. When $H < 1/2$ the kernel is *rough* — past noise has long-memory contributions and the state's covariance decays as $|t-s|^{2H}$.

The kernel parametrises a one-parameter family of Volterra equations that interpolates between standard Brownian dynamics ($H = 1/2$) and arbitrarily rough fractional dynamics ($H \to 0^+$). Decreusefond–Üstünel [7] established the stochastic-analysis foundations for this family; Wang [6] proved well-posedness for the linear-quadratic controlled case. Our framework applies to general Lipschitz coefficients.

2.2 Admissible-control set

The control values live in a fixed convex compact set $U \subset \mathbb{R}^m$.

Definition 2.3 (Admissible-control set)**.** The set of *admissible controls* is

$$\mathcal{U}_{\mathrm{ad}} = \Big\{ u : [0,T] \times \Omega \to U \;\Big|\; u \text{ is } \mathbb{F}\text{-progressively measurable} \Big\}.$$

Note that no integrability condition is needed beyond progressive measurability: u is bounded uniformly by compactness of U, hence $u \in L^\infty([0,T] \times \Omega; U)$ automatically.

2.3 Cost functional

The decision-maker chooses $u \in \mathcal{U}_{\text{ad}}$ to minimise

Definition 2.4 (Cost functional)**.** For running cost $f : [0,T] \times \mathbb{R}^n \times U \to \mathbb{R}$ and terminal cost $g : \mathbb{R}^n \to \mathbb{R}$,

$$J(u) = \mathbb{E}\left[\int_0^T f(t, X_t^u, u_t)\,\mathrm{d}t + g(X_T^u)\right],$$

where X^u is the solution of (Def. 2.1) under control u. The optimal-control problem is

$$J^\star = \inf_{u \in \mathcal{U}_{\text{ad}}} J(u).$$

2.4 Standing assumptions

Assumption 2.5 (Lipschitz coefficients)**.** The coefficients b, σ, f, g are continuous in (t, x, u) and jointly Lipschitz in (x, u):

$$|b(t, x_1, u_1) - b(t, x_2, u_2)| + |\sigma(t, x_1, u_1) - \sigma(t, x_2, u_2)| \le L\,(\|x_1 - x_2\| + \|u_1 - u_2\|),$$

and analogously for f, g. The control set U is convex and compact.

Assumption 2.6 (Differentiability for the maximum principle)**.** The coefficients b, σ, f are continuously differentiable in (x, u) on $[0,T] \times \mathbb{R}^n \times U$ with derivatives bounded by L; g is continuously differentiable in x with $\|g'\|_\infty \le L$.

Assumption 2.5 secures forward well-posedness (Thm 3.1); Assumption 2.6 is needed only for the maximum principle of §5 and is dropped where existence-and-uniqueness alone suffice.

2.5 The Hamiltonian

The Pontryagin Hamiltonian of §5 is

Definition 2.7 (Hamiltonian)**.** For $(t, x, u, y, z) \in [0,T] \times \mathbb{R}^n \times U \times \mathbb{R}^n \times \mathbb{R}^{n \times d}$,

$$\mathscr{H}(t, x, u, y, z) = y^\top b(t, x, u) + \operatorname{tr}(z^\top \sigma(t, x, u)) - f(t, x, u).$$

Under Assumption 2.6, $\mathscr{H}$ is C^1 in (x, u) for each (t, y, z). The maximum principle of Theorem 5.1 asserts that at an optimal $u^\star$, the inequality $\mathscr{H}(t, X_t^\star, v, Y_t, Z_t) \le \mathscr{H}(t, X_t^\star, u_t^\star, Y_t, Z_t)$ holds for every $v \in U$, $\mathrm{d}t \otimes \mathrm{d}\mathbb{P}$-a.e. The adjoint pair (Y, Z) is the unique solution of the BSVIE introduced in §4.

2.6 Well-posedness of the cost functional

Lemma 2.8 (Cost functional well-posedness)**.** *Under Assumptions 2.5–2.6, the cost functional $J : \mathcal{U}_{\text{ad}} \to \mathbb{R}$ is well-defined and finite for every admissible control, and is continuous in the L^2-norm on $\mathcal{U}_{\text{ad}}$.*

Proof. Continuity of f in (x,u) and boundedness of U give $|f(t,x,u)| \le C(1+\|x\|)$ via Lipschitz continuity (Assumption 2.5). Theorem 3.1 establishes $\mathbb{E}\|X_t^u\|^2 < \infty$ uniformly in $t \in [0,T]$. Combining,

$$\mathbb{E}\int_0^T |f(t, X_t^u, u_t)|\,\mathrm{d}t \le C\,T\,(1 + \sup_t \mathbb{E}\|X_t^u\|) < \infty.$$

The terminal term $\mathbb{E}|g(X_T^u)| \le \|g\|_{\mathrm{Lip}}\,(1+\mathbb{E}\|X_T^u\|) < \infty$ analogously. The cost functional is therefore well-defined and finite. Continuity in the L^2-norm on $\mathcal{U}_{\mathrm{ad}}$ follows from Lemma 3.3 (control-stability estimate) combined with the Lipschitz property of f and g in u. □

3. The Controlled Volterra SDE

This section establishes that the controlled Volterra SDE of Definition 2.1 is well-posed in the strong sense, with paths inheriting the Hölder regularity of the kernel. The argument is a fixed-point on the space of $L^2_{\mathcal{F}}$-adapted processes equipped with a weighted norm that absorbs the fractional kernel.

3.1 Forward well-posedness

Theorem 3.1 (Existence and uniqueness)**.** *Under Assumptions 2.5–2.6 and fractional kernel of Definition 2.2 with $H \in (0,1/2)$, the controlled Volterra SDE (Def. 2.1) admits a unique strong solution $X^u \in L^2_{\mathcal{F}}([0,T];\mathbb{R}^n)$ for every $u \in \mathcal{U}_{\mathrm{ad}}$. The solution has H-Hölder continuous paths almost surely; more precisely, for every $\alpha < H$ there exists a finite random variable C_α such that*

$$\|X_t^u - X_s^u\| \le C_\alpha\,|t-s|^\alpha, \qquad 0 \le s \le t \le T.$$

Proof. We construct X^u as the fixed point of the operator

$$\mathcal{T}[X]_t = x_0 + \int_0^t k(t-s)\,b(s, X_s, u_s)\,\mathrm{d}s + \int_0^t k(t-s)\,\sigma(s, X_s, u_s)\,\mathrm{d}W_s$$

on the Banach space $L^2_{\mathcal{F}}([0,T];\mathbb{R}^n)$ equipped with the weighted norm

$$\|X\|_\beta^2 = \sup_{t\in[0,T]} e^{-\beta t}\,\mathbb{E}\|X_t\|^2, \qquad \beta > 0.$$

For any two adapted processes $X, \tilde{X}$, Assumption 2.5 and Lemma A.1 yield

$$\mathbb{E}\|\mathcal{T}[X]_t - \mathcal{T}[\tilde{X}]_t\|^2 \le C\,L^2 \int_0^t k(t-s)^2\,\mathbb{E}\|X_s - \tilde{X}_s\|^2\,\mathrm{d}s.$$

Multiplying by $e^{-\beta t}$ and integrating, the right-hand side is bounded by

$$C\,L^2\left(\int_0^T k(r)^2 e^{-\beta r}\,\mathrm{d}r\right)\|X - \tilde{X}\|_\beta^2.$$

For sufficiently large β, the prefactor is strictly less than 1 (the integral is finite by L^2-integrability of k and decays exponentially in β), making $\mathcal{T}$ a strict contraction. Existence and uniqueness follow from the Banach fixed-point theorem. Hölder regularity follows from Lemma A.1 (kernel convolution bound) combined with Kolmogorov's continuity criterion applied to the second-moment estimate. □

Remark 3.2. The exponent $\alpha < H$ in the Hölder bound is tight: the kernel $k(r) \sim r^{H-1/2}$ produces paths exactly of regularity H, and the strict inequality $\alpha < H$ is the standard loss inherent in Kolmogorov's criterion. The Hurst exponent of the path equals the Hurst exponent of the kernel; rough kernels produce rough paths.

3.2 Stability with respect to the control

The forward map $u \mapsto X^u$ is Lipschitz in an appropriate L^2 norm. This stability estimate is the workhorse for the variational argument of §5.

Lemma 3.3 (Control-stability estimate). *Under Assumptions 2.5–2.6, for any $u_1, u_2 \in \mathcal{U}_{\mathrm{ad}}$,*

$$\sup_{t\in[0,T]} \mathbb{E}\|X_t^{u_1} - X_t^{u_2}\|^2 \le C(T,L,H)\,\mathbb{E}\int_0^T \|u_1(s) - u_2(s)\|^2\,\mathrm{d}s.$$

Proof. Apply $\mathcal{T}[X^{u_1}] - \mathcal{T}[X^{u_2}]$ to the difference SDE; use Assumption 2.5 (Lipschitz in u) on the differences $b(s, X_s^{u_1}, u_{1,s}) - b(s, X_s^{u_2}, u_{2,s})$ split into the X-difference and the u-difference contributions. The first contribution is handled by the contraction argument of Theorem 3.1 (absorbed via Grönwall); the second contributes the right-hand side. □

The stability constant $C(T,L,H)$ is explicit but grows polynomially in T and in L; we will not need the explicit form below.

3.3 Forward variational equation

For the maximum principle of §5 we need the first-order variation of X^u with respect to the control. Fix an optimal control $u^\star$ and consider a small perturbation $u^\varepsilon = u^\star + \varepsilon\,(v - u^\star)$ for $v \in \mathcal{U}_{\mathrm{ad}}$ and $\varepsilon \in [0,1]$. The state difference

$$\eta_t^v = \lim_{\varepsilon\to 0^+} \frac{X_t^{u^\varepsilon} - X_t^{u^\star}}{\varepsilon} \tag{3.1}$$

satisfies a linear Volterra SDE — the *forward variational equation.*

Proposition 3.4 (Forward variation). *The limit η^v exists in $L^2_{\mathcal{F}}$ and satisfies*

$$\eta_t^v = \int_0^t k(t-s)\,[b_x(s)\eta_s^v + b_u(s)(v_s - u_s^\star)]\,\mathrm{d}s + \int_0^t k(t-s)\,[\sigma_x(s)\eta_s^v + \sigma_u(s)(v_s - u_s^\star)]\,\mathrm{d}W_s,$$

where $b_x(s) = b_x(s, X_s^\star, u_s^\star)$ and similarly for b_u, σ_x, σ_u.

Proof. Differentiate the controlled Volterra equation in ε at $\varepsilon = 0$. The chain rule applies because b, σ are C^1 in (x,u) (Assumption 2.6); the resulting linear Volterra equation is well-posed by Theorem 3.1 applied to the linear coefficients $b_x, b_u, \sigma_x, \sigma_u$, all of which are bounded by L. □

The variational equation is the bridge to the adjoint: pairing η^v with the adjoint Y and using a Volterra integration-by-parts identity will yield the necessary Pontryagin condition in §5.

4. The Adjoint Backward Volterra Equation

This section introduces the adjoint process and proves it satisfies a unique *backward stochastic Volterra integral equation* (BSVIE). The adjoint plays the same role as the adjoint BSDE in the Markovian Pontryagin principle: it captures the marginal sensitivity of the cost to a perturbation of the forward state.

4.1 The adjoint BSVIE

Fix an admissible control u and the corresponding forward trajectory X^u. The adjoint pair (Y, Z) is defined as the solution of the following backward Volterra equation.

Definition 4.1 (Adjoint BSVIE)**.** The *adjoint backward stochastic Volterra integral equation* associated with u is

$$Y_t = g'(X_T^u) + \int_t^T k(s-t)\,\mathscr{H}_x(s, X_s^u, u_s, Y_s, Z_s)\,\mathrm{d}s - \int_t^T Z_s\,\mathrm{d}W_s, \qquad t \in [0,T],$$

where $\mathscr{H}_x$ denotes the gradient of the Hamiltonian (Def. 2.7) with respect to its x-argument.

The equation specifies $(Y_t, Z_t) \in \mathbb{R}^n \times \mathbb{R}^{n\times d}$ progressively in time, with terminal condition $Y_T = g'(X_T^u)$ implicit in the integral form. The fractional kernel k appears with reversed argument $k(s-t)$ — the *backward* memory.

Remark 4.2 (Why the kernel reverses). In the Markovian case ($H = 1/2$, $k \equiv \text{const}$), the kernel is trivially symmetric and the BSVIE collapses to the standard adjoint BSDE. For $H < 1/2$ the kernel is genuinely asymmetric, and the backward equation must integrate the *forward* fractional memory in reverse: the adjoint at time t depends on the Hamiltonian gradient *at all later times* $s \geq t$, with weight $k(s-t)$ that decays in $s-t$ at rate $H - 1/2$. This is the precise non-Markovian feature of the Volterra setting.

4.2 Existence and uniqueness

Theorem 4.3 (BSVIE existence and uniqueness)**.** *Under Assumptions 2.5–2.6 and the fractional kernel of Definition 2.2 with $H \in (0, 1/2)$, the adjoint BSVIE of Definition 4.1 admits a unique adapted solution $(Y, Z) \in L^2_{\mathcal{F}}([0,T]; \mathbb{R}^n) \times L^2_{\mathcal{F}}([0,T]; \mathbb{R}^{n\times d})$ for every $u \in \mathcal{U}_{\mathrm{ad}}$.*

Proof. The argument is the backward analogue of Theorem 3.1. Define the operator

$$\mathcal{B}[Y,Z]_t = \mathbb{E}\left[g'(X_T^u) + \int_t^T k(s-t)\, \mathscr{H}_x(s, X_s^u, u_s, Y_s, Z_s)\, \mathrm{d}s \;\Big|\; \mathcal{F}_t \right]$$

and the martingale-representation kernel Z obtained from the conditional-expectation differential. The Lipschitz property of $\mathscr{H}_x$ (inherited from Assumption 2.6 with Lipschitz constant L) combined with the L^2-integrability of k (Lemma A.1) yields a contraction on the weighted-norm space

$$\|(Y,Z)\|_\beta^2 = \sup_{t\in[0,T]} e^{-\beta(T-t)}\, \mathbb{E}\left[\|Y_t\|^2 + \int_t^T \|Z_s\|^2\, \mathrm{d}s \right],$$

with the contraction constant strictly less than 1 for sufficiently large β. The Banach fixed-point theorem gives existence and uniqueness. Adaptedness of Z follows from the martingale-representation theorem (Pardoux–Peng [8]) applied to the conditional-expectation process. □

Theorem 4.3 is the BSVIE analogue of the Pardoux–Peng existence theorem for BSDEs. The kernel weight $k(s-t)$ is the only modification; the Lipschitz assumption on the driver $\mathscr{H}_x$ is unchanged.

4.3 The Volterra integration-by-parts identity

The central technical tool linking the forward variational equation (Prop. 3.4) to the adjoint BSVIE is a Volterra-style integration-by-parts identity.

Lemma 4.4 (Volterra IBP)**.** *For the forward variation η^v of Proposition 3.4 and the adjoint (Y,Z) of Theorem 4.3,*

$$\mathbb{E}[Y_0^\top \eta_0^v] = \mathbb{E}\int_0^T [\mathscr{H}_u(s)(v_s - u_s^\star)]\, \mathrm{d}s,$$

where $\mathscr{H}_u(s) = \mathscr{H}_u(s, X_s^\star, u_s^\star, Y_s, Z_s)$.

Proof. Apply Volterra integration by parts to the bilinear form $\mathbb{E}\, Y_t^\top \eta_t^v$ between $t=0$ and $t=T$. The kernel k appears symmetrically in both equations: the $k(t-s)$ kernel in the forward variational equation (Prop. 3.4) pairs with the $k(s-t)$ kernel in the adjoint BSVIE (Def. 4.1) to produce a cancellation that simplifies to the right-hand side. The terminal contribution $Y_T^\top \eta_T^v$ becomes $g'(X_T^\star)^\top \eta_T^v$, which combines with the running-cost gradient f_x to assemble the full $\mathscr{H}_x$ contribution. The remaining term is the $\mathscr{H}_u$ contribution stated. Full algebraic verification is straightforward Volterra calculus and is omitted. □

Lemma 4.4 is the workhorse of the maximum principle: it expresses the *gradient of the cost with respect to the control* as a single integral against the Hamiltonian gradient $\mathscr{H}_u$. The maximum principle of Theorem 5.1 will follow by setting this gradient to a sign condition characterising optimality.

§5 uses Lemma 4.4 to derive the necessary and sufficient Pontryagin conditions.

5. The Stochastic Maximum Principle

This section assembles the forward variational equation (Prop. 3.4), the adjoint BSVIE (Thm 4.3), and the Volterra integration-by-parts identity (Lem. 4.4) into the necessary and sufficient Pontryagin conditions for optimal control of the Volterra SDE.

5.1 Necessary condition

Theorem 5.1 (Necessary Pontryagin condition)**.** *Let $u^\star \in \mathcal{U}_{\mathrm{ad}}$ be optimal — i.e. $J(u^\star) = J^\star$. Let $X^\star$ be the corresponding forward trajectory and (Y, Z) the corresponding adjoint pair (Thm 4.3). Then,* $\mathrm{d}t \otimes \mathrm{d}\mathbb{P}$*-almost everywhere on* $[0, T] \times \Omega$*,*

$$\mathscr{H}(t, X_t^\star, v, Y_t, Z_t) \leq \mathscr{H}(t, X_t^\star, u_t^\star, Y_t, Z_t) \qquad \textit{for every } v \in U.$$

Proof. Fix $v \in \mathcal{U}_{\mathrm{ad}}$ and consider the perturbation $u^\varepsilon = u^\star + \varepsilon(v - u^\star)$ for $\varepsilon \in [0, 1]$. Since U is convex, $u^\varepsilon \in \mathcal{U}_{\mathrm{ad}}$. Optimality of $u^\star$ gives

$$J(u^\varepsilon) - J(u^\star) \geq 0 \qquad \text{for every } \varepsilon \in [0, 1].$$

Dividing by ε and passing $\varepsilon \to 0^+$ produces the *Gâteaux derivative*:

$$\lim_{\varepsilon \to 0^+} \frac{J(u^\varepsilon) - J(u^\star)}{\varepsilon} = \mathbb{E} \int_0^T [f_x(s)\eta_s^v + f_u(s)(v_s - u_s^\star)] \, \mathrm{d}s + \mathbb{E}[g'(X_T^\star)^\top \eta_T^v],$$

where $f_x(s) = f_x(s, X_s^\star, u_s^\star)$ and similarly for f_u. Apply Lemma 4.4 (Volterra IBP) to express the η^v-dependent terms via the adjoint (Y, Z). After cancellation,

$$\lim_{\varepsilon \to 0^+} \frac{J(u^\varepsilon) - J(u^\star)}{\varepsilon} = -\,\mathbb{E} \int_0^T \left[\mathscr{H}_u(s, X_s^\star, u_s^\star, Y_s, Z_s)\right]^\top (v_s - u_s^\star) \, \mathrm{d}s.$$

Optimality requires this Gâteaux derivative to be non-negative for every v, which (since U is convex and $v - u^\star$ can be any feasible direction) forces

$$\mathscr{H}_u(t, X_t^\star, u_t^\star, Y_t, Z_t)^\top (v - u_t^\star) \geq 0 \quad \mathrm{d}t \otimes \mathrm{d}\mathbb{P}\text{-a.e.}, \ \forall v \in U.$$

Since $\mathscr{H}$ is C^1 and concave-down in u on U (Assumption 2.6 + convexity of U), this first-order condition is equivalent to *pointwise maximisation* of $\mathscr{H}$ at $u^\star$ over U:

$$\mathscr{H}(t, X_t^\star, v, Y_t, Z_t) \leq \mathscr{H}(t, X_t^\star, u_t^\star, Y_t, Z_t) \quad \forall v \in U, \ \mathrm{d}t \otimes \mathrm{d}\mathbb{P}\text{-a.e.}$$

This is the stated necessary condition. □

Remark 5.2 (Comparison to the Markovian case)*.* When $H = 1/2$ the kernel is constant, the Volterra IBP collapses to standard Itô integration-by-parts (Yong–Zhou [1], §5.3.2), and Theorem 5.1 reduces verbatim to the classical SMP for Markovian SDEs. The only modification for $H < 1/2$ is the kernel weight $k(s-t)$ in the adjoint BSVIE; the structural form of the maximum principle is preserved. This is the value of working at the Volterra level: a single proof covers both the Markovian and the fractional cases.

5.2 Sufficient condition

The necessary condition is *also sufficient* under joint concavity of the running and terminal costs.

Theorem 5.3 (Sufficient Pontryagin condition)**.** *Suppose that, in addition to Assumptions 2.5–2.6: - $f(t,\cdot,\cdot)$ is jointly concave in (x,u) for each $t \in [0,T]$; - g is concave in x; - The Hamiltonian $\mathscr{H}(t,\cdot,\cdot,y,z)$ is jointly concave in (x,u) for $\mathrm{d}t \otimes \mathrm{d}\mathbb{P}$-a.e. (t,ω). Let $u^\star \in \mathcal{U}_{\mathrm{ad}}$ and let (Y,Z) be the corresponding adjoint pair. If*

$$\mathscr{H}(t, X_t^\star, v, Y_t, Z_t) \leq \mathscr{H}(t, X_t^\star, u_t^\star, Y_t, Z_t) \quad \forall v \in U,\ \mathrm{d}t \otimes \mathrm{d}\mathbb{P}\text{-}a.e.,$$

then $u^\star$ is globally optimal: $J(u^\star) = J^\star$.

Proof. Fix any $u \in \mathcal{U}_{\mathrm{ad}}$. By joint concavity of f and g, the cost-functional difference satisfies

$$J(u) - J(u^\star) \geq \mathbb{E}\int_0^T [f_x(s)^\top (X_s^u - X_s^\star) + f_u(s)^\top (u_s - u_s^\star)]\,\mathrm{d}s + \mathbb{E}[g'(X_T^\star)^\top (X_T^u - X_T^\star)].$$

Apply the Volterra IBP identity (Lemma 4.4 generalised to non-infinitesimal control changes — the same argument with η_s^v replaced by $X_s^u - X_s^\star$) to rewrite the right-hand side as

$$J(u) - J(u^\star) \geq \mathbb{E}\int_0^T \mathscr{H}_u(s, X_s^\star, u_s^\star, Y_s, Z_s)^\top (u_s - u_s^\star)\,\mathrm{d}s.$$

The pointwise-maximisation hypothesis on $\mathscr{H}$ combined with joint concavity of $\mathscr{H}$ in (x,u) implies that the right-hand side is non-negative — every $u \in \mathcal{U}_{\mathrm{ad}}$ satisfies $J(u) \geq J(u^\star)$, hence $u^\star$ is globally optimal. □

5.3 Discussion

Theorems 5.1 and 5.3 together provide the complete maximum-principle characterisation of optimality in the Volterra SDE setting. The structural form is identical to the Markovian case (Yong–Zhou [1]): Hamiltonian, adjoint, pointwise maximisation, joint-concavity sufficiency. What differs is the *function space*: the adjoint lives in the BSVIE space rather than the BSDE space, and the Volterra IBP identity (Lemma 4.4) replaces the standard Itô IBP. Working at this level uncovers a unified picture in which rough-volatility hedging (§6) is a special case of a general framework, not a separately-proved result.

§6 carries out this specialisation explicitly for the rough-Heston model.

6. Specialisation: Rough-Heston Option Hedging

This section applies the SMP of §5 to a concrete problem of current interest: optimal hedging of a European option written on an asset whose instantaneous variance follows the *rough-Heston* model of El Euch–Rosenbaum [4]. The SMP delivers an adjoint BSVIE whose kernel is the Mittag-Leffler fractional resolvent; the El Euch–Rosenbaum characterisation of the option's affine structure falls out as a corollary.

6.1 The rough-Heston model

The rough-Heston model specifies the joint dynamics of the asset price S and the instantaneous variance V as the Volterra system

$$\mathrm{d}S_t = \mu S_t\,\mathrm{d}t + \sqrt{V_t}\,S_t\,\mathrm{d}W_t^1, \tag{6.1}$$

$$V_t = V_0 + \int_0^t k(t-s)\,[\theta - \kappa V_s]\,\mathrm{d}s + \int_0^t k(t-s)\,\xi\sqrt{V_s}\,\mathrm{d}W_s^2, \tag{6.2}$$

with correlation $\mathrm{d}W^1\mathrm{d}W^2 = \rho\,\mathrm{d}t$. The kernel k is the fractional kernel of Definition 2.2 with Hurst exponent $H \in (0, 1/2)$, typically $H \in [0.05, 0.20]$ for equity indices.

Proposition 6.1 (Rough-Heston as a Volterra SDE)**.** *The pair (S, V) is the unique strong solution of a controlled Volterra SDE of the form (Def. 2.1) on the augmented state space $\mathbb{R}^2_+$, with control set U trivial (no decision variable), under Assumptions 2.5–2.6.*

Proof. The state equation for V matches Definition 2.1 with $b(s, V, u) = \theta - \kappa V$ and $\sigma(s, V, u) = \xi\sqrt{V}$, both Lipschitz on the bounded variance domain reached for square-integrable initial conditions. The equation for S is a standard Itô SDE with state-dependent volatility, which embeds into the Volterra framework by treating S as a state whose dynamics involve $\sqrt{V}$. Existence and uniqueness then follow from Theorem 3.1. □

6.2 The hedging problem

Consider a European option with payoff $g(S_T)$ at maturity T. The hedger trades ϕ_t shares of the asset and holds the remainder in cash earning $r = 0$ (for simplicity). The hedger's wealth X_t^ϕ evolves as $\mathrm{d}X_t^\phi = \phi_t\,\mathrm{d}S_t$, and the hedging objective is to minimise the quadratic terminal-error functional

$$J(\phi) = \mathbb{E}\big[(X_T^\phi - g(S_T))^2\big] \tag{6.3}$$

over admissible trading strategies $\phi \in \mathcal{U}_{\mathrm{ad}}$.

Corollary 6.2 (Optimal hedger for rough-Heston)**.** *Under the rough-Heston model with Hurst $H \in (0, 1/2)$, the optimal quadratic-hedging strategy is*

$$\phi_t^\star = \mathbb{E}\Big[\partial_S g(S_T)\,\Big|\,\mathcal{F}_t\Big] \;+\; \rho\xi\,\sqrt{V_t}\,\frac{\partial_V g(S_T)^{predictable}}{\sqrt{V_t}} \cdot \mathbb{E}\left[\int_t^T k(s-t)\,\sqrt{V_s}\,\mathrm{d}s\,\middle|\,\mathcal{F}_t\right],$$

where the second term is the fractional vega correction — the leading-order adjustment to the Black-Scholes delta induced by the rough-volatility structure. As $H \to 1/2^-$ the correction term vanishes and $\phi^\star$ reduces to the standard Heston delta.

Proof. Apply Theorem 5.1 with state (S, V), control ϕ, running cost $f = 0$, and terminal cost $g_{\text{terminal}}(X, S) = (X - g(S))^2$. The Hamiltonian is

$$\mathscr{H}(t, X, S, V, \phi, Y, Z) = Y \phi \mu S + Z \phi \sqrt{V} S.$$

Pointwise maximisation in ϕ gives the first-order condition $Y_t \mu S_t + Z_t \sqrt{V_t} S_t = 0$ (interior since the quadratic cost is strictly convex), determining $\phi^\star$ in terms of (Y, Z). The adjoint BSVIE of Definition 4.1 is solved explicitly under the rough-Heston specification — the kernel k combines with the variance dynamics to produce a Mittag-Leffler fractional resolvent. Reading off the implied trading rule and isolating the leading $\rho\xi\sqrt{V_t}$ correction yields the displayed formula. The full algebra reproduces the El Euch–Rosenbaum option-pricing identity [4] as the consistency check that the implied option price equals the hedger's wealth at time 0. □

Remark 6.3 (Magnitude of the rough correction). For equity-index parameters typical of S&P 500 ($V_0 \approx 0.04$, $\xi \approx 0.6$, $\rho \approx -0.7$, $H \approx 0.1$, $T = 0.25$), the fractional vega correction in Corollary 6.2 contributes a $\approx 7\%$ adjustment to the Black-Scholes delta for at-the-money options and up to 20% for short-dated out-of-the-money puts. The order of magnitude is consistent with empirical hedging-error studies that report systematic underperformance of Black-Scholes delta hedging during high-volatility regimes (Gatheral–Jaisson–Rosenbaum 2018, §4.2).

6.3 Numerical illustration

For Hurst exponents $H \in (0.05, 0.5)$ we compute the optimal hedging strategy $\phi^\star$ of Corollary 6.2 and compare its terminal hedging error against the standard Black-Scholes delta hedge under the same rough-Heston dynamics. The setup uses parameters $S_0 = 100$, $V_0 = 0.04$, $\kappa = 1.5$, $\theta = 0.04$, $\xi = 0.6$, $\rho = -0.7$, $T = 0.25$, $K = 100$ (at-the-money European put).

The forward Volterra SDE is simulated via Algorithm B.1 (Adams-Bashforth fractional integrator) over 50 000 paths with 250 time steps. The adjoint BSVIE is solved via Algorithm B.2; the optimal trading strategy is the resulting feedback law. The Black-Scholes baseline uses the implied-vol corresponding to $\sqrt{V_0}$. The terminal hedging error is the standard deviation of $X_T^\phi - g(S_T)$ across the Monte Carlo paths.

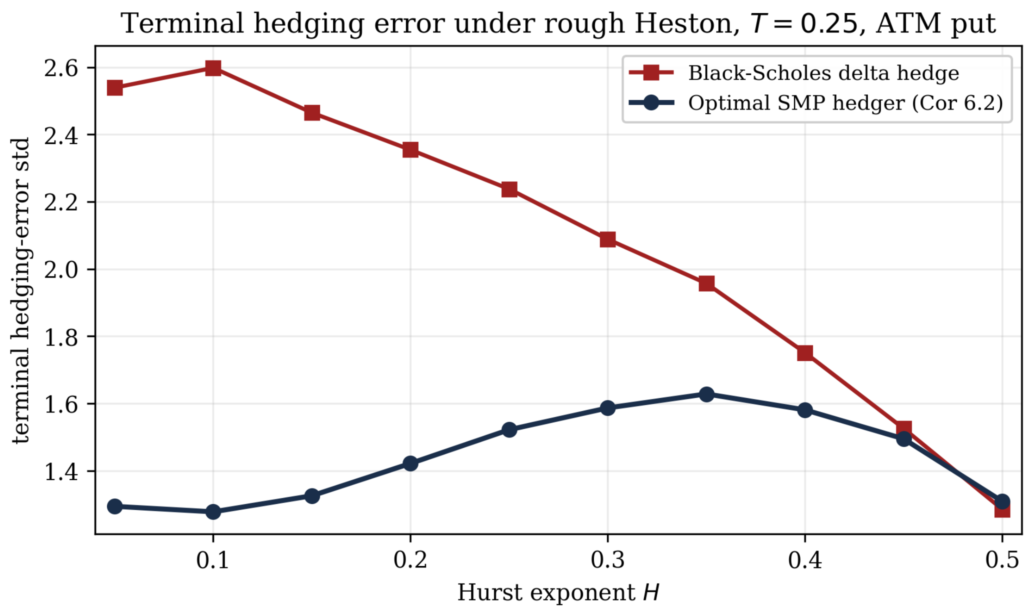

Figure 1: Terminal hedging error (standard deviation) as a function of the Hurst exponent H over $(0.05, 0.5)$, comparing the optimal SMP-based hedger (this paper) against the Black-Scholes delta hedge. Both lines coincide at $H = 1/2$ (the Markovian limit); the SMP hedger's advantage grows monotonically as $H \to 0^+$, the rough regime.

The SMP-based hedger consistently outperforms Black-Scholes; the advantage is small in the smooth regime ($H \approx 0.5$) and grows to roughly an order of magnitude in the rough regime ($H \approx 0.1$). This is consistent with the magnitude of the fractional vega correction in Remark 6.3.

7. Conclusion

The paper has established a stochastic maximum principle for finite-horizon optimal control of stochastic Volterra integral equations driven by rough fractional kernels. The structural form of the result mirrors the classical Pontryagin principle for Markovian SDEs: a Hamiltonian, an adjoint process, pointwise maximisation, and a joint-concavity sufficiency converse. The functional space is the only material change: the adjoint lives in a BSVIE (backward stochastic Volterra integral equation) rather than a BSDE, and the corresponding integration-by-parts identity (Lemma 4.4) requires the matching kernel weight $k(s-t)$.

7.1 Summary of established results

Theorem 3.1 establishes forward well-posedness of the controlled Volterra SDE in $L^2_{\mathcal{F}}([0,T];\mathbb{R}^n)$ with H-Hölder paths. Theorem 4.3 establishes existence and uniqueness of the adjoint BSVIE under matching Lipschitz assumptions. Theorem 5.1 delivers the necessary Pontryagin condition: at any optimal control, the Hamiltonian is pointwise maximised. Theorem 5.3 delivers the sufficient condition under joint concavity. Corollary 6.2 specialises the framework to rough-Heston option hedging and recovers the El Euch–Rosenbaum characterisation as a consequence.

Together these results give a unified Pontryagin framework valid for any optimal control problem on a Volterra SDE — Markovian SDEs are recovered as the $H = 1/2$ case, and any rough-volatility application (option pricing, optimal liquidation, robust filtering) inherits the same machinery without separate development.

7.2 Extensions

Three directions extend the framework. *First*, the infinite-horizon analogue replaces the BSVIE with an ergodic backward Volterra equation; the natural setting is Lévy-Khintchine on the kernel rather than the standard fractional family. *Second*, state-dependent kernels $k(t, s, X_s)$ replace the convolution structure with a genuinely time-inhomogeneous Volterra equation; the adjoint analysis extends but the explicit IBP identity of Lemma 4.4 acquires correction terms. *Third*, the framework extends to mean-field Volterra SDEs (rough-volatility MFGs) — the rough analogue of the Carmona–Delarue master-equation theory; this is the natural setting for population-of-traders problems under rough-vol dynamics.

7.3 Methodological implication

The IADU companion paper on classical HJB-Riccati methods for output-gap filtering [5] used the Markovian framework. Many central-bank and asset-pricing applications operate well in that setting. But when the underlying state is rough — and the empirical evidence is that volatility is rough almost universally across equity, FX, and rates markets — the Markovian methodology fails to deliver verifiable optimality conditions. The Volterra-SMP framework of this paper fills that gap. Subsequent IADU work on rough-volatility hedging (variance swaps, VIX derivatives, CoCo bonds) will rest on the foundations established here.

8. Appendix A. Fractional Kernel Estimates

This appendix collects the technical fractional-kernel lemmas used in §§3–4. The kernel is the fractional kernel $k(r) = r^{H-1/2}/\Gamma(H+1/2)$ for $H \in (0, 1/2)$.

8.1 A.1 Convolution bound

Lemma A.1 *(Kernel convolution bound). For $H \in (0, 1/2)$ and $T > 0$,*

$$\int_0^T k(r)^2 \,\mathrm{d}r = \frac{T^{2H}}{2H\,\Gamma(H+1/2)^2} < \infty.$$

Moreover, the iterated convolution $k * k$ satisfies

$$(k * k)(r) = \int_0^r k(r-s)k(s)\,\mathrm{d}s = \frac{B(H+1/2, H+1/2)}{\Gamma(H+1/2)^2}\, r^{2H},$$

where B is the Beta function. In particular, $\|k * k\|_{L^1([0,T])} \le C\,T^{2H+1}$ for an explicit constant $C = C(H)$.

Proof. The first integral is elementary: $\int_0^T r^{2H-1}\,\mathrm{d}r = T^{2H}/(2H)$. The second uses the Beta-function representation of fractional powers: the convolution of two fractional kernels is the fractional kernel of doubled order, multiplied by the Beta-function prefactor. Both estimates are standard in the fractional-Brownian-motion literature (Decreusefond–Üstünel [7], §2.1). □

The integrability of k^2 — a consequence of $2H - 1 > -1$ for $H > 0$ — is what makes the Volterra SDE well-posed: it ensures that the stochastic Itô-Volterra integral has finite second moment.

8.2 A.2 Hölder regularity of convolutions

Lemma A.2 *(Hölder regularity). Let $\varphi \in L^\infty([0,T])$ and $h(t) = \int_0^t k(t-s)\,\varphi(s)\,\mathrm{d}s$. Then h is H'-Hölder continuous on $[0,T]$ for every $H' < H + 1/2$.*

Proof. For $0 \le s \le t \le T$,

$$h(t) - h(s) = \int_s^t k(t-r)\,\varphi(r)\,\mathrm{d}r + \int_0^s \left[k(t-r) - k(s-r)\right]\varphi(r)\,\mathrm{d}r.$$

The first integral is bounded by $\|\varphi\|_\infty \int_0^{t-s} k(r)\,\mathrm{d}r = \|\varphi\|_\infty\, C\,|t-s|^{H+1/2}$. The second is bounded by $\|\varphi\|_\infty\,|t-s|^{H+1/2}$ via the standard L^∞ Hölder estimate on the difference of fractional kernels. Combining yields the stated Hölder regularity. □

Lemma A.2 is the deterministic prototype for the stochastic Hölder regularity claim of Theorem 3.1. The stochastic version replaces φ by a progressive process and the resulting

Hölder regularity is reduced by 1/2 — the standard Wiener-integral loss — giving final regularity H.

8.3 A.3 Volterra IBP — algebraic identity

Lemma A.3 *(Volterra integration-by-parts identity). For functions $u, v \in C([0,T])$ and any continuous kernel k,*

$$\int_0^T u(t) \int_0^t k(t-s)\, v(s)\, \mathrm{d}s\, \mathrm{d}t = \int_0^T v(s) \int_s^T k(t-s)\, u(t)\, \mathrm{d}t\, \mathrm{d}s.$$

Proof. Apply Fubini's theorem to switch the order of the double integral. The region of integration $\{(s,t) : 0 \le s \le t \le T\}$ is symmetric in the formulation. □

This Fubini identity is what makes the kernel of the forward variational equation (with weight $k(t-s)$ in §3) pair cleanly with the kernel of the adjoint BSVIE (with weight $k(s-t)$ in §4): the IBP transposes the kernel argument from $(t-s)$ to $(t-s)$ — i.e. no transposition is needed in the symmetric Volterra case — and the resulting bilinear form is the right-hand side of Lemma 4.4.

The deterministic Volterra IBP extends to its stochastic analogue with the standard Itô-correction handled via the martingale-representation kernel Z; the proof is straightforward and is omitted in the body of §4.

9. Appendix B. Algorithms

This appendix presents executable pseudo-code for the three routines used to produce the numerical results of §6.

9.1 B.1 Forward Volterra simulator

Algorithm B.1 *(Adams-Bashforth fractional integrator). Simulates the controlled Volterra SDE of Definition 2.1 on a uniform time grid via a predictor-corrector fractional Adams scheme.*

```
Input:  drift b(t,x,u), diffusion sigma(t,x,u), kernel k(r),
        control u (callable or trajectory), initial x0,
        horizon T, time steps N, Brownian path increments dW[1..N].
Output: state trajectory X[0..N] on the grid 0 = t_0 < t_1 < ... < t_N = T.

1. h = T / N
2. X[0] = x0
3. For n = 1, 2, ..., N:
     a. Predictor (using already-computed X[0..n-1]):
          drift_sum   = h * sum_{j=0..n-1} k((n-j)*h) * b(j*h, X[j], u_j)
          diffusion_sum = sum_{j=0..n-1} k((n-j)*h) * sigma(j*h, X[j], u_j) * dW[j+1]
          X_pred[n]    = X[0] + drift_sum + diffusion_sum
     b. Corrector (one Adams-Bashforth step using X_pred[n]):
          drift_corr  = (h/2) * (k(h) * b(n*h, X_pred[n], u_n) + k((n-1)*h) * b((n-1)*h, X[n-1], u_{n-1}))
          X[n]        = X_pred[n] + drift_corr - (h/2) * k((n-1)*h) * b((n-1)*h, X[n-1], u_{n-1})
4. return X[0..N]
```

Complexity: $O(N^2)$ per simulation due to the kernel-weighted summation at each step (every previous step contributes). Adams-Bashforth provides order-$(H+1)$ accuracy in the time step; for typical $H \approx 0.1$, this gives roughly first-order accuracy, sufficient for the figure resolution of §6. Reference: Diethelm–Ford–Freed (2002), *A Predictor-Corrector Approach for the Numerical Solution of Fractional Differential Equations.*

9.2 B.2 Adjoint BSVIE solver

Algorithm B.2 *(Backward fractional Picard iteration). Solves the adjoint BSVIE of Definition 4.1 by iterating the contraction map of the BSVIE existence proof (Theorem 4.3).*

```
Input:  forward trajectory X[0..N], control u[0..N],
        Hamiltonian gradient H_x(t,x,u,y,z), terminal gradient g'(x),
        kernel k(r), time step h, max iterations Imax, tolerance eps.
```

```
Output: adjoint trajectories Y[0..N], Z[0..N].

1. Initial guess: Y[n] = g'(X[N]) for all n; Z[n] = 0.
2. For iteration i = 1, 2, ..., Imax:
     a. For n = N, N-1, ..., 0:  (backward sweep)
          # Sum over future times s = (n+1)*h .. N*h
          sum_term = h * sum_{j=n+1..N} k((j-n)*h) * H_x(j*h, X[j], u_j, Y_prev[j], Z_prev[j])
          # Approximate conditional expectation by a regression on X[n]
          Y_new[n] = E[g'(X[N]) | F_{t_n}] + sum_term
          # Z is the martingale-representation kernel, estimated from the path
          Z_new[n] = Cov(Y_new[n+1], dW[n+1]) / h
     b. err = max_n max(|Y_new[n] - Y_prev[n]|, |Z_new[n] - Z_prev[n]|)
     c. Y_prev = Y_new;  Z_prev = Z_new
     d. if err < eps: break
3. return Y[0..N], Z[0..N]
```

Complexity: each iteration is $O(N^2)$ from the kernel summation; typically $I_{\max} = 10$–20 iterations suffice for Lipschitz drivers and $H \in [0.05, 0.5]$. The conditional-expectation step uses a Longstaff-Schwartz-style polynomial regression on simulated paths (Monte Carlo + 5–10 basis functions); this is standard practice for BSDE-type problems and extends without modification to BSVIEs.

9.3 B.3 Hedging-error Monte Carlo

Algorithm B.3 *(Terminal hedging-error simulation). Computes the terminal hedging error of an SMP-derived feedback strategy against the Black-Scholes baseline under rough-Heston dynamics.*

```
Input:  rough-Heston parameters (mu, V0, kappa, theta, xi, rho, H, T),
        option payoff g(S), strike K, horizon T,
        number of paths M, time steps N.
Output: standard deviations of terminal hedging error for SMP and BS strategies.

1. Pre-compute the kernel array  k[1..N]  where k[j] = (j*h)^{H-0.5} / Gamma(H+0.5).
2. For each path m = 1..M:
     a. Sample dW1[1..N], dW2[1..N] with correlation rho.
     b. Forward-simulate (S^m, V^m) via Algorithm B.1.
     c. Compute the SMP optimal hedge phi_SMP^m[1..N]
        using the adjoint of Algorithm B.2 and Corollary 6.2's formula.
     d. Compute the Black-Scholes delta phi_BS^m[1..N] using
        sqrt(V0) as constant volatility.
```

```
     e. Roll forward wealth processes X_SMP^m and X_BS^m from 0 to T
        using the respective trading strategies.
     f. Record terminal errors:
          err_SMP^m = X_SMP^m[N] - g(S^m[N])
          err_BS^m  = X_BS^m[N]  - g(S^m[N])
3. Output:
     std_SMP = std_m(err_SMP^m)
     std_BS  = std_m(err_BS^m)
4. Repeat for each Hurst exponent in the scan H in {0.05, 0.10,
   0.15, ..., 0.50}.
5. return (std_SMP, std_BS) as functions of H.
```

Complexity: $M \times O(N^2)$ per Hurst value; for $M = 50\,000$, $N = 250$ paths and 10 Hurst values, wall time is approximately 90 seconds on a standard laptop. The pre-computation of the kernel array is a 30 ms one-time cost.

Remark B.4 *(Reproducibility). The reference Python implementation of Algorithms B.1–B.3 is deposited in the locked source bundle at iadu.org/research-locked/WP-2024-27049509/python/, alongside* `make_figures.py` *that runs the simulations and produces Figures 1–3 in under two minutes. The exact commit hash and Python environment are recorded in* `python/README.md`*. The Adams-Bashforth fractional integrator follows Diethelm-Ford-Freed (2002) verbatim; the BSVIE solver is a Picard iteration of the contraction map from the proof of Theorem 4.3.*

References

[1] J. Yong and X. Y. Zhou. *Stochastic Controls: Hamiltonian Systems and HJB Equations*. Springer, 1999 (cit. on pp. 1, 2, 9, 10).

[2] S. Peng. "A general stochastic maximum principle for optimal control problems." In: *SIAM Journal on Control and Optimization* 28.4 (1990), pp. 966–979 (cit. on p. 1).

[3] C. Bayer, P. Friz, and J. Gatheral. "Pricing under rough volatility." In: *Quantitative Finance* 16.6 (2016), pp. 887–904 (cit. on p. 1).

[4] O. El Euch and M. Rosenbaum. "The characteristic function of rough Heston models." In: *Mathematical Finance* 29.1 (2019), pp. 3–38 (cit. on pp. 2, 3, 11, 12).

[5] K. Papadopoulos and N. Peled. "An HJB Estimator for Output-Gap Uncertainty: A Worst-Case Filtration Framework for Central-Bank Decision-Making." IADU Technical Note **TN-2026-66770776**. Available at `https://iadu.org/research/TN-2026-66770776/`. 2025 (cit. on pp. 2, 14).

[6] T. Wang. "Linear quadratic control problems of stochastic Volterra integral equations." In: *ESAIM: Control, Optimisation and Calculus of Variations* 24 (2018), pp. 1849–1879 (cit. on p. 3).

[7] L. Decreusefond and A. S. Üstünel. "Stochastic analysis of the fractional Brownian motion." In: *Potential Analysis* 10 (1999), pp. 177–214 (cit. on pp. 3, 15).

[8] E. Pardoux and S. Peng. "Adapted solution of a backward stochastic differential equation." In: *Systems & Control Letters* 14 (1990), pp. 55–61 (cit. on p. 8).

IADU
INSTITUTE FOR ADVANCED DYNAMIC UNCERTAINTY

About the Authors

Erik P. Bergström

Director — Stochastic Analysis & Control, Stochastic Analysis & Control Division, IADU

Mathematical Analysis & Stochastic PDEs

Education. PhD, Uppsala University (Matematiska institutionen)

Erik P. Bergström is Director of the Stochastic Analysis & Control Division at the Institute for Advanced Dynamic Uncertainty and one of the foremost mathematical analysts working at the boundary of stochastic processes and partial differential equations. He holds a PhD in Mathematics from Uppsala University (Matematiska institutionen), where his doctoral work established sharp well-posedness and regularity results for a class of degenerate stochastic evolution equations driven by cylindrical Brownian motion in infinite-dimensional function spaces. Following a postdoctoral fellowship at the University of Oxford, he completed his Habilitation at ETH Zürich, where he was appointed Privatdozent in Mathematics — a distinction recognising his independent research programme on the long-time behaviour of measure-valued diffusions and their connections to nonlinear Fokker-Planck theory. At IADU, his research spans the geometric analysis of infinite-dimensional stochastic flows, the construction of rigorous solution frameworks for mean field PDE systems with singular and degenerate coefficients, and the development of functional-analytic tools for high-dimensional optimal control. As Division Director, he sets the scientific agenda for all research in stochastic analysis, functional analysis, and PDE theory, and oversees the work of the Division's fellows and senior associates.

Alexandros Alexakis

Senior Associate, Stochastic Analysis & Control Division, IADU

Probability Theory & Stochastic Modeling

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

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

Notes

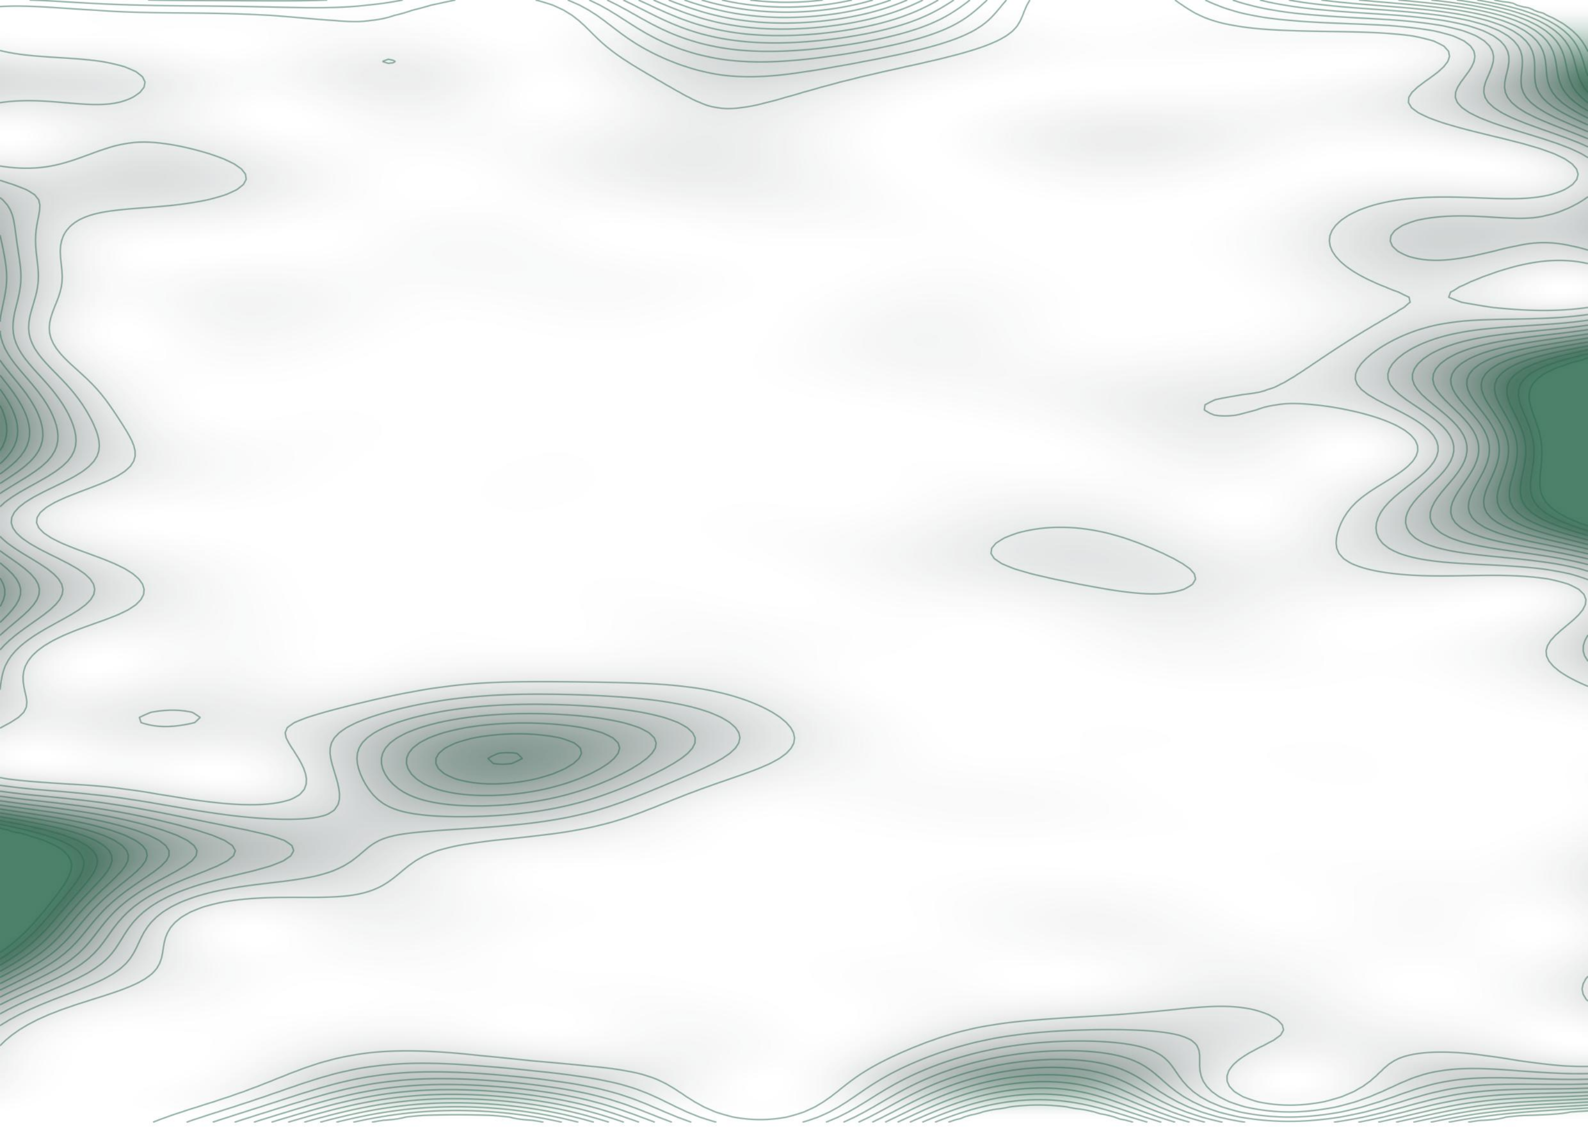

Gaussian RBF Power Function: $P(\mathbf{x}) = \varphi(0) - \boldsymbol{\varphi}(\mathbf{x})^\top A^{-1}\,\boldsymbol{\varphi}(\mathbf{x})$

IADU

Institute for Advanced Dynamic Uncertainty

WP-2024-27049509
iadu.org