Weierstrass Institute for
Applied Analysis and Stochastics

Leibniz Association

Primal and dual optimal stopping with signatures

Christian Bayer

Joint work with: L. Pelizzari, J. Schoenmakers

ICCF 2024

Outline

Recent trend for using processes with memory in finance and beyond:

- Rough volatility: Model stochastic volatility by fractional Brownian motion, e.g., the rough Bergomi model:

$$dS_t = \sqrt{v_t} S_t dZ_t,$$
$$v_t = \xi(t) \exp\left(\eta \widehat{W}_t - \frac{1}{2}\eta^2 t^{2H}\right), \ \widehat{W}_t := \int_0^t K(t-s) dW_s, \ K(r) := \sqrt{2H} r^{H-\frac{1}{2}}.$$

- Order flow models by self-exciting jump processes, e.g., Hawkes processes.
- Statistical mechanics models based on Generalized Langevin Equations.

Recent trend for using processes with memory in finance and beyond:

- Rough volatility: Model stochastic volatility by fractional Brownian motion, e.g., the rough Bergomi model:

$$dS_t = \sqrt{v_t} S_t dZ_t,$$
$$v_t = \xi(t) \exp\left(\eta \widehat{W}_t - \frac{1}{2}\eta^2 t^{2H}\right), \; \widehat{W}_t := \int_0^t K(t-s) dW_s, \; K(r) := \sqrt{2H} r^{H-\frac{1}{2}}.$$

- Order flow models by self-exciting jump processes, e.g., Hawkes processes.
- Statistical mechanics models based on Generalized Langevin Equations.

Many numerical methods rely on the Markov property: (pricing) PDEs, polynomial regression methods, dynamic programming, ...

Value function

$$v(t,x) := \sup_{\tau \in \mathcal{S},\, \tau \geq t} \mathbb{E}\left[Y_{\tau \wedge T} \mid X_t = x \right]$$

- $X_t \in \mathbb{R}^d$ denotes an underlying Markov (asset price + additional factors) process, $d \geq 1$
- Y_t denotes the (discounted) cash-flow process, e.g., $Y_t = g(X_t)$.
- $\mathbb{E}$ w.r.t. a pricing measure $\mathbb{P}$, $\mathcal{S}$ the set of $(\mathcal{F}_t)$-stopping times, $(\mathcal{F}_t)$ generated by X.

Value function

$$v(t,x) := \sup_{\tau \in \mathcal{S},\, \tau \geq t} \mathbb{E}\left[Y_{\tau \wedge T} \mid X_t = x\right]$$

- $X_t \in \mathbb{R}^d$ denotes an underlying Markov (asset price + additional factors) process, $d \geq 1$
- Y_t denotes the (discounted) cash-flow process, e.g., $Y_t = g(X_t)$.
- $\mathbb{E}$ w.r.t. a pricing measure $\mathbb{P}$, $\mathcal{S}$ the set of $(\mathcal{F}_t)$-stopping times, $(\mathcal{F}_t)$ generated by X.

Dynamic programming principle

$$v(t,x) \approx \max(\mathbb{E}[v(t+\Delta, X_{t+\Delta}) \mid X_t = x], g(x))$$

- Approximate $\mathbb{E}[v(t+\Delta, X_{t+\Delta}) \mid X_t = x]$ by regression based on family of basis functions $\mathcal{A}$, e.g., $\mathcal{A} = \mathrm{Pol}_{\deg \leq n}(\mathbb{R}^d)$, $\frac{1}{M}\sum_{i=1}^{M}\left(\overline{v}(t+\Delta, X^{(i)}_{t+\Delta}) - \sum_{\phi \in \mathcal{A}} c_\phi \phi(X^{(i)}_t)\right)^2 \xrightarrow{c_\phi} \min!$
- Curse of dimensionality of the functional approximation problem.

$$\sup_{\tau \in \mathcal{S}} \mathbb{E}\left[Y_{\tau \wedge T}\right] = \inf_{M \in \mathcal{M}_0} \mathbb{E}\left[\sup_{t \in [0,T]} (Y_t - M_t)\right]$$

- $\mathcal{M}_0$ denotes martingales M with $M_0 = 0$, and $\mathbb{E}\left[||M||_\infty\right] < \infty$.
- An optimizer M^* is given by the martingale of the Doob–Meyer decomposition of the Snell envelop $V_t := \operatorname{ess\,sup}_{\tau \in \mathcal{S}_t} \mathbb{E}\left[Y_{\tau \wedge T} \mid \mathcal{F}_t\right]$.

Dual martingale approach [Rogers ’02, Haugh and Kogan ’04]

$$\sup_{\tau \in \mathcal{S}} \mathbb{E}\left[Y_{\tau\wedge T}\right] = \inf_{M \in \mathcal{M}_0} \mathbb{E}\left[\sup_{t\in[0,T]} (Y_t - M_t)\right]$$

- $\mathcal{M}_0$ denotes martingales M with $M_0 = 0$, and $\mathbb{E}\left[\|M\|_\infty\right] < \infty$.
- An optimizer M^* is given by the martingale of the Doob–Meyer decomposition of the Snell envelop $V_t := \operatorname{ess\,sup}_{\tau\in\mathcal{S}_t} \mathbb{E}\left[Y_{\tau\wedge T} \mid \mathcal{F}_t\right]$.
- If, for simplicity, X is a diffusion process driven by a Brownian motion W, then $V_t = v(t, X_t)$, and the Doob-Meyer decomposition reads

$$M_t^* = \int_0^t \partial_x v(s, X_s)\sigma(X_s)\,\mathrm{d}W_s. \quad \mathrm{d}X_t = \mu(X_t)\,\mathrm{d}t + \sigma(X_t)\,\mathrm{d}W_t.$$

This motivates the *ansatz* $M_t = \int_0^t f_\theta(s, X_s)\,\mathrm{d}W_s$, for f_θ in some suitable parametric function space, $\theta \in \Theta$.

The signature $\widehat{\mathbb{X}}^{<\infty}_{0,t}$, i.e., the collection of all iterated integrals, determines the path $X|_{[0,t]}$.
Hence, the process $t \mapsto \widehat{\mathbb{X}}^{<\infty}_{0,t}$ is a Markov process, even when X is not.

Dynamic programming

- Markovian ansatz:

$$\mathbb{E}[V_{t+\Delta t} \mid \mathcal{F}_t] = f_\theta(t, X_t)$$

- Non-Markovian ansatz:

$$\mathbb{E}[V_{t+\Delta t} \mid \mathcal{F}_t] = f(t, X|_{[0,t]}) = f_\theta\left(\widehat{\mathbb{X}}^{\leq N}_{0,t}\right)$$

Dual martingale approach

- Markovian ansatz:

$$M_t = \int_0^t f_\theta(s, X_s)\, \mathrm{d}W_s$$

- Non-Markovian ansatz:

$$M_t = \int_0^t f\left(s, X|_{[0,s]}\right) \mathrm{d}W_s = \int_0^t f_\theta\left(\widehat{\mathbb{X}}^{\leq N}_{0,s}\right) \mathrm{d}W_s$$

Controlled differential equation

Let $x : [0,T] \to \mathbb{R}^d$ be a smooth path, $V : \mathbb{R}^e \to \mathbb{R}^{e \times d}$ smooth, $y_0 \in \mathbb{R}^e$, and consider

$$\mathrm{d}y(t) = V(y(t))\,\mathrm{d}x(t), \quad t \in [0,T], \quad y(0) = y_0.$$

Controlled differential equation

Let $x : [0,T] \to \mathbb{R}^d$ be a smooth path, $V : \mathbb{R}^e \to \mathbb{R}^{e\times d}$ smooth, $y_0 \in \mathbb{R}^e$, and consider

$$\mathrm{d}y(t) = V(y(t))\,\mathrm{d}x(t), \quad t \in [0,T], \quad y(0) = y_0.$$

- y solves an ODE: $\dot{y}(t) = V(y(t))\dot{x}(t)$, but difficult to generalize to rough paths.

Controlled differential equation

Let $x : [0, T] \to \mathbb{R}^d$ be a smooth path, $V : \mathbb{R}^e \to \mathbb{R}^{e\times d}$ smooth, $y_0 \in \mathbb{R}^e$, and consider

$$dy(t) = V(y(t))\,dx(t), \quad t \in [0, T], \quad y(0) = y_0.$$

- y solves an ODE: $\dot{y}(t) = V(y(t))\dot{x}(t)$, but difficult to generalize to rough paths.
- First order expansion: For $s < u < t$, $y(u) = y(s) +$ H.O.T., implying that

  $V(y(u)) = V(y(s))+$H.O.T., and hence $y(t) = y(s) + V(y(s))x_{s,t} +$ H.O.T., $x_{s,t} \coloneqq x(t) - x(s)$.

Controlled differential equation

Let $x : [0,T] \to \mathbb{R}^d$ be a smooth path, $V : \mathbb{R}^e \to \mathbb{R}^{e\times d}$ smooth, $y_0 \in \mathbb{R}^e$, and consider

$$\mathrm{d}y(t) = V(y(t))\,\mathrm{d}x(t), \quad t \in [0,T], \quad y(0) = y_0.$$

- y solves an ODE: $\dot{y}(t) = V(y(t))\dot{x}(t)$, but difficult to generalize to rough paths.
- First order expansion: For $s < u < t$, $y(u) = y(s) + \text{H.O.T.}$, implying that

$$V(y(u)) = V(y(s)) + \text{H.O.T.}, \text{ and hence } y(t) = y(s) + V(y(s))x_{s,t} + \text{H.O.T.}, \quad x_{s,t} := x(t) - x(s).$$

- Second order expansion: $y(u) = y(s) + V(y(s))x_{s,u} + \text{H.O.T.}$, implying that

$$V(y(u)) = V(y(s)) + DV(y(s))V(y(s))x_{s,u}, \quad y(t) = y(s) + V(y(s))x_{s,t} + DV(y(s))V(y(s))\mathbb{x}_{s,t} + \text{H.O.T.}$$

$$\mathbb{x}^{\mathtt{ij}}_{s,t} := \int_s^t x^i_{s,u}\mathrm{d}x^j(u) = \int_{s<t_1<t_2<t} \mathrm{d}x^i(t_1)\,\mathrm{d}x^j(t_2), \quad i,j = 1,\dots,d.$$

Controlled differential equation

Let $x : [0, T] \to \mathbb{R}^d$ be a smooth path, $V : \mathbb{R}^e \to \mathbb{R}^{e\times d}$ smooth, $y_0 \in \mathbb{R}^e$, and consider

$$\mathrm{d}y(t) = V(y(t))\,\mathrm{d}x(t), \quad t \in [0, T], \quad y(0) = y_0.$$

- y solves an ODE: $\dot{y}(t) = V(y(t))\dot{x}(t)$, but difficult to generalize to rough paths.
- First order expansion: For $s < u < t$, $y(u) = y(s) +$ H.O.T., implying that

  $V(y(u)) = V(y(s))+$H.O.T., and hence $y(t) = y(s) + V(y(s))x_{s,t} +$ H.O.T., $x_{s,t} := x(t) - x(s)$.
- Second order expansion: $y(u) = y(s) + V(y(s))x_{s,u} +$ H.O.T., implying that

$V(y(u)) = V(y(s)) + DV(y(s))V(y(s))x_{s,u}$, $y(t) = y(s) + V(y(s))x_{s,t} + DV(y(s))V(y(s))\mathbb{x}_{s,t} +$ H.O.T.

$$\mathbb{x}^{\mathtt{ij}}_{s,t} := \int_s^t x^i_{s,u}\mathrm{d}x^j(u) = \int_{s<t_1<t_2<t} \mathrm{d}x^i(t_1)\,\mathrm{d}x^j(t_2), \; i, j = 1, \dots, d.$$

- Third order expansion: involves iterated integrals of order three...

- Given a (possibly random, but for now smooth) path $X : [0,T] \to \mathbb{R}^d$.
- W.l.o.g., $X(0) = 0$. Denote $\widehat{X}(t) := (t, X(t))$, $X^0(t) := t$.

Signature

The signature is the collection of all iterated integrals,

$$\widehat{\mathbb{X}}^{\mathtt{i}_1 \cdots \mathtt{i}_n}_{s,t} := \int_{s<t_1<\cdots<t_n<t} d\widehat{X}^{\mathtt{i}_1}_{t_1} \cdots d\widehat{X}^{\mathtt{i}_n}_{t_n}, \quad \mathtt{i}_1, \ldots, \mathtt{i}_n \in \{\mathtt{0}, \ldots, \mathtt{d}\}.$$

- Given a (possibly random, but for now smooth) path $X : [0,T] \to \mathbb{R}^d$.
- W.l.o.g., $X(0) = 0$. Denote $\widehat{X}(t) := (t, X(t))$, $X^0(t) := t$.

Signature

The signature is the collection of all iterated integrals,

$$\widehat{\mathbb{X}}^{\mathtt{i}_1 \cdots \mathtt{i}_n}_{s,t} := \int_{s<t_1<\cdots<t_n<t} d\widehat{X}^{\mathtt{i}_1}_{t_1} \cdots d\widehat{X}^{\mathtt{i}_n}_{t_n}, \quad \mathtt{i}_1, \ldots, \mathtt{i}_n \in \{\mathtt{0}, \ldots, \mathtt{d}\}.$$

- Natural structure: Associate $\widehat{\mathbb{X}}^{\mathtt{i}_1 \cdots \mathtt{i}_n}_{s,t}$ with the multi-index $\mathtt{i}_1 \cdots \mathtt{i}_n$.
  Operations on multi-indices: addition, scalar multiplication, concatenation product.
- Motivation: Natural relations between iterated integrals, e.g. $\widehat{\mathbb{X}}^{\mathtt{ij}}_{s,t} + \widehat{\mathbb{X}}^{\mathtt{ji}}_{s,t} = \widehat{\mathbb{X}}^{\mathtt{i}}_{s,t} \widehat{\mathbb{X}}^{\mathtt{j}}_{s,t}$.
- Obtain formal power series in $1 + d$ non-commutating variables $\mathtt{0}, \ldots, \mathtt{d}$.

- The concatenation product on $\operatorname{span}\{\mathtt{0},\dots,\mathtt{d}\}$ is equivalent to the tensor product $\otimes$ on $(\mathbb{R}^{1+d})^{\otimes k}$.
- The signature is formally defined as an element of the (extended) tensor algebra $T((\mathbb{R}^{1+d}))$ (with product $\otimes$), i.e.,

$$\widehat{\mathbb{X}}^{<\infty}_{s,t} := \sum_{n=0}^{\infty} \sum_{\mathtt{i}_1,\dots,\mathtt{i}_n \in \{\mathtt{0},\dots,\mathtt{d}\}} \widehat{X}^{\mathtt{i}_1\cdots\mathtt{i}_n}_{s,t} e_{\mathtt{i}_1} \otimes \cdots \otimes e_{\mathtt{i}_n} \in T((\mathbb{R}^{1+d})) := \prod_{n=0}^{\infty} (\mathbb{R}^{1+d})^{\otimes n}.$$

  $\widehat{\mathbb{X}}^{\leq N}_{s,t} \in T^N(\mathbb{R}^{1+d})$ denotes the truncation to level N, the projection to level equal to n is denoted by $\pi_n : T^N(\mathbb{R}^{1+d}) \to (\mathbb{R}^{1+d})^{\otimes n}$.
- Let $\mathcal{W}_{1+d}$ denote the linear span of words $\mathtt{w}$ in the letters $\{\mathtt{0},\mathtt{1},\dots,\mathtt{d}\}$. Bracket defined for $\mathtt{w} = \mathtt{i}_1\cdots\mathtt{i}_k$, $\mathbf{a} = \sum_{n=0}^{\infty} \sum_{\mathtt{i}_1,\dots,\mathtt{i}_n \in \{\mathtt{0},\mathtt{1},\dots,\mathtt{d}\}} a^{\mathtt{i}_1\cdots\mathtt{i}_n} e_{\mathtt{i}_1} \otimes \cdots \otimes e_{\mathtt{i}_n} \in T((\mathbb{R}^d))$ by $\langle \mathtt{w}, \mathbf{a}\rangle := a^{\mathtt{w}}$.

Chen's rule

$$\widehat{\mathbb{X}}^{<\infty}_{s,u} \otimes \widehat{\mathbb{X}}^{<\infty}_{u,t} = \widehat{\mathbb{X}}^{<\infty}_{s,t}, \quad 0 \le s \le u \le t \le T.$$

Chen's rule

$$\widehat{\mathbb{X}}^{<\infty}_{s,u} \otimes \widehat{\mathbb{X}}^{<\infty}_{u,t} = \widehat{\mathbb{X}}^{<\infty}_{s,t}, \quad 0 \le s \le u \le t \le T.$$

- Shuffle product on $\mathcal{W}_d$: For words $\mathtt{w}, \mathtt{v}$ and letters $\mathtt{i}, \mathtt{j}$ defined by

$$\mathtt{w} ⧢ \varnothing := \varnothing ⧢ \mathtt{w} := \mathtt{w}, \quad \mathtt{wi} ⧢ \mathtt{vj} := (\mathtt{w} ⧢ \mathtt{vj})\mathtt{i} + (\mathtt{wi} ⧢ \mathtt{vj})\mathtt{j}.$$

- Example: $\mathtt{12} ⧢ \mathtt{34} = \mathtt{1234} + \mathtt{1324} + \mathtt{1342} + \mathtt{3124} + \mathtt{3142} + \mathtt{3412}$.

Shuffle identity

$$\forall \ell_1, \ell_2 \in \mathcal{W}_{1+d} : \left\langle \ell_1, \widehat{\mathbb{X}}^{<\infty}_{s,t} \right\rangle \left\langle \ell_2, \widehat{\mathbb{X}}^{<\infty}_{s,t} \right\rangle = \left\langle \ell_1 ⧢ \ell_2, \widehat{\mathbb{X}}^{<\infty}_{s,t} \right\rangle.$$

Chen's rule

$$\widehat{\mathbb{X}}^{<\infty}_{s,u} \otimes \widehat{\mathbb{X}}^{<\infty}_{u,t} = \widehat{\mathbb{X}}^{<\infty}_{s,t}, \quad 0 \le s \le u \le t \le T.$$

- Shuffle product on $\mathcal{W}_d$: For words $\mathtt{w}, \mathtt{v}$ and letters $\mathtt{i}, \mathtt{j}$ defined by

$$\mathtt{w} \sqcup\!\sqcup \varnothing := \varnothing \sqcup\!\sqcup \mathtt{w} := \mathtt{w}, \quad \mathtt{wi} \sqcup\!\sqcup \mathtt{vj} := (\mathtt{w} \sqcup\!\sqcup \mathtt{vj})\mathtt{i} + (\mathtt{wi} \sqcup\!\sqcup \mathtt{vj})\mathtt{j}.$$

- Example: $\mathtt{12} \sqcup\!\sqcup \mathtt{34} = \mathtt{1234} + \mathtt{1324} + \mathtt{1342} + \mathtt{3124} + \mathtt{3142} + \mathtt{3412}$.

Shuffle identity

$$\forall \ell_1, \ell_2 \in \mathcal{W}_{1+d} : \left\langle \ell_1, \widehat{\mathbb{X}}^{<\infty}_{s,t} \right\rangle \left\langle \ell_2, \widehat{\mathbb{X}}^{<\infty}_{s,t} \right\rangle = \left\langle \ell_1 \sqcup\!\sqcup \ell_2, \widehat{\mathbb{X}}^{<\infty}_{s,t} \right\rangle.$$

Path encoding

$\widehat{\mathbb{X}}^{<\infty}_{0,T}$ determines the path $X|_{[0,T]}$. Note: this holds due to time extension $\widehat{X}$.

- For $\widehat{\mathbb{X}} : \Delta_T \to T^N(\mathbb{R}^{1+d})$, $\Delta_T := \{ (s,t) \mid 0 \le s \le t \le T \}$, let

$$\left\|\widehat{\mathbb{X}}\right\|_\alpha := \max_{n=1,\dots,N} \left(\sup_{0\le s<t\le T} \frac{\left|\pi_n(\widehat{\mathbb{X}}_{s,t})\right|}{|t-s|^{n\alpha}} \right)^{1/n}.$$

Rough paths

Given $\alpha \in]0,1[$, the set of (geometric) α-Hölder rough paths is the closure of $\left\{ \widehat{\mathbb{X}}^{\le\lfloor 1/\alpha \rfloor}_{\cdot,\cdot} \,\middle|\, X \text{ smooth} \right\}$ under $\|\cdot\|_\alpha$. It is denoted by $\mathscr{C}^\alpha_g([0,T];\mathbb{R}^{1+d})$.

- Given a rough path $\widehat{\mathbb{X}}$, we can construct $\widehat{\mathbb{X}}^{<\infty}$ in a unique, pathwise, continuous way.
- Example: Let W be a Brownian motion, set $\mathbb{W}(\omega) : \Delta_T \to T^2(\mathbb{R}^d)$ by

$$W^i_{s,t} := W^i_t - W^i_s, \quad W^{i,j}_{s,t} := \int_s^t (W^i_u - W^i_s) \circ \mathrm{d}W^j_u, \quad 1 \le i,j \le d.$$

This a.s. defines a rough path for $1/3 < \alpha < 1/2$, i.e., $\mathbb{W} \in \mathscr{C}^\alpha_g$ a.s.

Continuous functionals $f : \widehat{\mathscr{C}}_g^{\alpha}([0,T]) \to \mathbb{R}$ can be approximated by linear functionals $\widehat{\mathbb{X}} \mapsto \left\langle \ell, \widehat{\mathbb{X}}_{0,T}^{<\infty} \right\rangle$, $\ell \in \mathcal{W}_{1+d}$.

- This is a consequence of Stone–Weierstrass and the shuffle identity (and holds on compact subsets of $\widehat{\mathscr{C}}_g^{\alpha}([0,T];\mathbb{R}^{1+d})$).

Continuous functionals $f : \widehat{\mathscr{C}}^{\alpha}_{g}([0,T]) \to \mathbb{R}$ can be approximated by linear functionals $\widehat{\mathbb{X}} \mapsto \left\langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,T} \right\rangle$, $\ell \in \mathcal{W}_{1+d}$.

- This is a consequence of Stone–Weierstrass and the shuffle identity (and holds on compact subsets of $\widehat{\mathscr{C}}^{\alpha}_{g}([0,T];\mathbb{R}^{1+d})$).

For every rough stochastic process $\widehat{\mathbb{X}}$, the process $t \mapsto \widehat{\mathbb{X}}^{<\infty}_{0,t}$ is a Markov process.

- Every rough path $\mathbb{X}$ *with one strictly monotone component* is uniquely determined by its signature.
- Assuming that X_0 is deterministic and, hence, $\mathcal{F}_0$ is trivial, the above result follows.

[Becker, Cheredito, Jentzen '19] consider the problem $\sup_{0 \leq \tau \leq 1} \mathbb{E}\left[W_\tau^H\right]$,

where W^H is fractional Brownian motion with Hurst index H – connection to rough stochastic volatility models.

[Becker, Cheredito, Jentzen '19] consider the problem $\sup_{0\le\tau\le1} \mathbb{E}\left[W_\tau^H\right]$,

where W^H is fractional Brownian motion with Hurst index H – connection to rough stochastic volatility models.

- Fix a time-grid $0 = t_0 < t_1 < \cdots < t_J = 1$, and define a Markov process $X_j \in \mathbb{R}^J$ by

$$
\begin{aligned}
X_0 &= (0, 0, \dots, 0) \\
X_1 &= (W_{t_1}^H, 0, \dots, 0) \\
X_2 &= (W_{t_1}^H, W_{t_2}^H, 0, \dots, 0) \\
&\vdots
\end{aligned}
$$

- Use deep neural networks to parameterize stopping decisions $f_j(X_j) \approx \mathrm{DNN}_j(X_j; \theta)$ – "stop at time j unless stopped earlier".

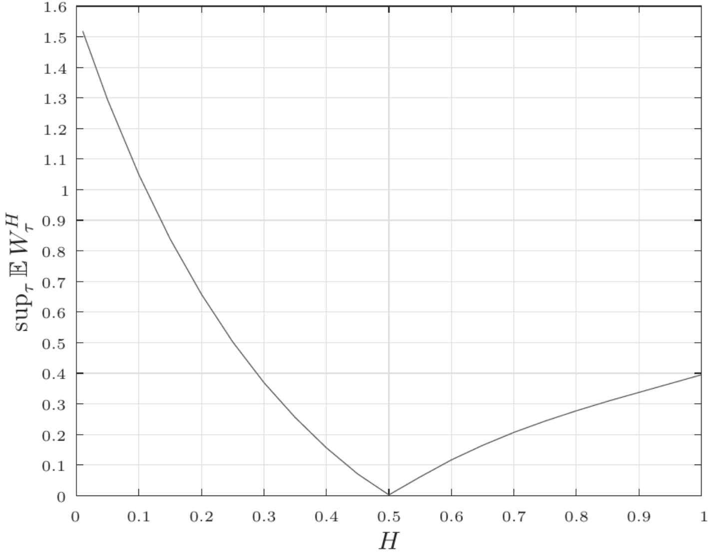

Figure: Plot from [Becker, Cheridito, Jentzen '19].

On a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ we are given:

- A stochastic process $(X_t)_{t\in[0,T]}$ such that $\widehat{X}_t := (t, X_t)$ extends to an α-Hölder rough path $\widehat{\mathbb{X}}$, $X_0 \equiv 0$. Alternatively, we consider a random variable taking values in $\widehat{\mathscr{C}}_g^{\alpha}([0,T];\mathbb{R}^{1+d})$.
- A continuous reward-process $(Y_t)_{t\in[0,T]}$ adapted to the filtration $(\mathcal{F}_t)_{t\in[0,T]}$ generated by $\widehat{\mathbb{X}}$ such that $\mathbb{E}\left[\|Y\|_{\infty;[0,T]}\right] < \infty$.

Optimal stopping problem

Let $\mathcal{S}$ be the set of $(\mathcal{F}_t)_{t\in[0,T]}$-stopping times taking values in $[0,T]$. Solve

$$\sup_{\tau\in\mathcal{S}} \mathbb{E} Y_\tau .$$

- Could also consider more general stochastic optimal control problems.

Definition (Space of stopped rough paths)

For $\alpha \in]0,1[$, $T > 0$, define $\Lambda_T^\alpha := \bigsqcup_{t\in[0,T]} \widehat{\mathscr{C}}_g^\alpha([0,T];\mathbb{R}^{1+d})$ equipped with the final topology of

$$\phi : [0,T] \times \widehat{\mathscr{C}}_g^\alpha([0,T];\mathbb{R}^{1+d}) \to \Lambda_T^\alpha, \quad \phi(t,\widehat{\mathbb{x}}) = \widehat{\mathbb{x}}|_{[0,t]}.$$

- Λ_T^α is a Polish space with metric $d(\widehat{\mathbb{x}}|_{[0,t]}, \widehat{\mathbb{y}}|_{[0,s]}) := \left\| \widehat{\mathbb{x}} - \widetilde{\mathbb{y}} \right\|_{\alpha;[0,t]} + |t-s|$ for $s \le t$, where $\widetilde{\mathbb{y}}$ is a piecewise constant extension (up to the 0-component) of $\widehat{\mathbb{y}}|_{[0,s]}$ to $[0,t]$.

Lemma

For any progressively measurable process Z, there is a measurable function $f : (\Lambda_T^\alpha, \mathcal{B}(\Lambda_T^\alpha)) \to \mathbb{R}$ s.t. $\forall t \in [0,T] : Z_t = f(\widehat{\mathbb{X}}|_{[0,t]})$ a.s.

Following [Kalsi, Lyons, Perez Arribas ’20], a method of solving stochastic optimal control problems using signatures can be described as follows:

1. Controls u_t are continuous functions of the path $\phi(\widehat{X}|_{[0,t]})$ and, hence, of the signature $\theta(\widehat{\mathbb{X}}^{<\infty}_{0,t})$ – and similarly for the loss function.
2. We may approximate $\theta(\widehat{\mathbb{X}}^{<\infty}_{0,T})$ by linear functionals $\left\langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,T} \right\rangle$.
3. Interchange expectation and truncate the signature at level N.
4. Optimize $\ell \mapsto \left\langle \ell, \mathbb{E}\left[\widehat{\mathbb{X}}^{\leq N}_{0,T}\right]\right\rangle$.

Pathwise density for steps **1.** + **2.** with high probability is proved in [Kalsi, Lyons, Perez Arribas ’20] for a data-driven optimal execution problem.

Given a tensor normalization λ (in the sense of [Chevyrev–Oberhauser '22]), consider

$$L^{\lambda}_{\text{sig}}(\Lambda^{\alpha}_{T}) := \left\{ \Lambda^{\alpha}_{T} \ni \widehat{\mathbb{X}}|_{[0,t]} \mapsto \left\langle \ell,\ \lambda(\widehat{\mathbb{X}}^{<\infty}_{0,t}) \right\rangle \,\middle|\, \ell \in \mathcal{W}_{1+d} \right\} \subset C_b\left(\Lambda^{\alpha}_{T}; \mathbb{R}\right).$$

Lemma

Let μ be a finite measure on Λ^{α}_{T} such that there is $\beta > \alpha$ with $\mu\left(\Lambda^{\alpha}_{T} \setminus \Lambda^{\beta}_{T}\right) = 0$. Then for every $f \in L^p(\Lambda^{\alpha}_{T}, \mu)$, $1 \le p < \infty$, there is $f_n \in L^{\lambda}_{\text{sig}}(\Lambda^{\alpha}_{T})$ s.t.

$$\|f - f_n\|_{L^p(\Lambda^{\alpha}_{T};\mu)} \xrightarrow[n\to\infty]{} 0.$$

- Similar approach by [Cuchiero, Schmocker, Teichmann '23] using weighted spaces.

- Let ν be a probability measure on $\widehat{\mathscr{C}}_g^\alpha([0,T];\mathbb{R}^{1+d})$ s.t. there is $\beta > \alpha$ with ν being supported on $\widehat{\mathscr{C}}_g^\beta([0,T];\mathbb{R}^{1+d}) \subset \widehat{\mathscr{C}}_g^\alpha([0,T];\mathbb{R}^{1+d})$.
- Example: the Wiener measure with $\alpha < 1/2$ and any $\alpha < \beta < 1/2$.
- Take μ to be the push-forward of $\mathrm{d}t \otimes \nu$ under $\phi : [0,T] \times \widehat{\mathscr{C}}_g^\alpha([0,T];\mathbb{R}^{1+d}) \ni (t, \widehat{\mathbb{x}}) \mapsto \widehat{\mathbb{x}}|_{[0,t]} \in \Lambda_T^\alpha$.
- Then, for $p = 2$, there are $\ell^n \in \mathcal{W}_{1+d}$ s.t.

$$\int_0^T \mathbb{E}\left[\left(f(\widehat{\mathbb{X}}|_{[0,t]}) - \left\langle \ell^n, \lambda(\widehat{\mathbb{X}}_{0,t}^{<\infty})\right\rangle\right)^2\right] \mathrm{d}t \to 0.$$

Consider the discrete time optimal stopping problem (Bermudan option): Given $\mathcal{T} := \{ 0 = t_0 < \cdots < t_K = T \}$, let

$$V_0^{\mathcal{T}} := \sup_{\tau \in \mathcal{S}_0^{\mathcal{T}}} \mathbb{E}[Y_\tau], \quad V_{t_k}^{\mathcal{T}} := \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_k^{\mathcal{T}}} \mathbb{E}[Y_\tau \mid \mathcal{F}_{t_k}],$$

where $\mathcal{S}_k^{\mathcal{T}}$ denotes the set of $(\mathcal{F}_t)_{t \in \mathcal{T}}$-stopping times taking values in $\{ t_k, \dots, t_K \}$.

Longstaff–Schwartz algorithm (2001), general version

An optimal stopping time is given by τ_0^*, where τ_k^* is recursively defined by $\tau_K^* := T$ and

$$\tau_k^* := t_k \mathbb{1}_{\left\{ Y_{t_k} \geq \mathbb{E}[Y_{\tau_{k+1}^*} | \mathcal{F}_{t_k}] \right\}} + \tau_{k+1}^* \mathbb{1}_{\left\{ Y_{t_k} < \mathbb{E}[Y_{\tau_{k+1}^*} | \mathcal{F}_{t_k}] \right\}}.$$

Optimality follows from dynamic programming $V_{t_k}^{\mathcal{T}} = \max\left(Y_{t_k}, \mathbb{E}[V_{t_{k+1}}^{\mathcal{T}} \mid \mathcal{F}_{t_k}]\right)$ together with the fact that $\min\{ t_m \in \{ t_k, \dots, t_K \} : V_{t_m}^{\mathcal{T}} = Y_{t_m} \}$ is optimal.

Let $\widetilde{\tau}_k = \widetilde{\tau}_k^{N,\Delta t,M}$ denote a sequence of stopping times starting at $\widetilde{\tau}_K = T$ and

$$\widetilde{\tau}_k := t_k \mathbb{1}_{\{Y_{t_k} \geq \psi_k^{N,\Delta t,M}(\widehat{\mathbb{X}}|_{[0,t_k]})\}} + \widetilde{\tau}_{k+1} \mathbb{1}_{\{Y_{t_k} < \psi_k^{N,\Delta t,M}(\widehat{\mathbb{X}}|_{[0,t_k]})\}}, \quad k \geq 1, \text{ where}$$

$\psi_k^{N,\Delta t,M}(\widehat{\mathbb{X}}|_{[0,t_k]}) = \left\langle \ell_k^{N,\Delta t,M}, \lambda(\widetilde{\mathbb{X}}_{0,t_k}^{\leq N}) \right\rangle$ with $\widetilde{\mathbb{X}}_{0,t_k}^{\leq N}$ denoting an approximation of the truncated signature based on a grid with mesh size Δt, and

$$\ell_k^{N,\Delta t,M} := \underset{\ell \in \mathcal{W}_{1+d}^{\leq N}}{\arg\min} \frac{1}{M} \sum_{m=1}^{M} \left(Y_{\widetilde{\tau}_{k+1}}^{(m)} - \left\langle \ell, \lambda(\widetilde{\mathbb{X}}_{0,t_k}^{\leq N,(m)}) \right\rangle \right)^2 .$$

Let $\widetilde{\tau}_k = \widetilde{\tau}_k^{N,\Delta t,M}$ denote a sequence of stopping times starting at $\widetilde{\tau}_K = T$ and

$$\widetilde{\tau}_k := t_k \mathbb{1}_{\{Y_{t_k} \geq \psi_k^{N,\Delta t,M}(\widehat{\mathbb{X}}|_{[0,t_k]})\}} + \widetilde{\tau}_{k+1} \mathbb{1}_{\{Y_{t_k} < \psi_k^{N,\Delta t,M}(\widehat{\mathbb{X}}|_{[0,t_k]})\}}, \quad k \geq 1, \text{ where}$$

$\psi_k^{N,\Delta t,M}(\widehat{\mathbb{X}}|_{[0,t_k]}) = \left\langle \ell_k^{N,\Delta t,M}, \lambda(\widetilde{\mathbb{X}}_{0,t_k}^{\leq N}) \right\rangle$ with $\widetilde{\mathbb{X}}_{0,t_k}^{\leq N}$ denoting an approximation of the truncated signature based on a grid with mesh size Δt, and

$$\ell_k^{N,\Delta t,M} := \underset{\ell \in \mathcal{W}_{1+d}^{\leq N}}{\arg\min} \frac{1}{M} \sum_{m=1}^{M} \left(Y_{\widetilde{\tau}_{k+1}}^{(m)} - \left\langle \ell, \lambda(\widetilde{\mathbb{X}}_{0,t_k}^{\leq N,(m)}) \right\rangle \right)^2.$$

Theorem

Assuming $\mathbb{E}\left[\|Y\|_{\infty;\mathcal{T}}\right] < \infty$, *and a technical support condition, we have that*

$$\widetilde{V}_0^{\mathcal{T}} := \max\left(Y_0, \frac{1}{M}\sum_{m=1}^{M} Y_{\widetilde{\tau}_1}^{(m)}\right) \xrightarrow[\Delta t \to 0]{M,N\to\infty} V_0^{\mathcal{T}} \text{ a.s.}$$

- Using independent samples for $\widetilde{V}_0^{\mathcal{T}}$ gives a low-biased estimator, i.e., $\mathbb{E}\left[\widetilde{V}_0^{\mathcal{T}}\right] \leq V_0^{\mathcal{T}}$.
- The optimal stopping problem is formulated in terms of the filtration $(\mathcal{F}_{t_k})_{k=0}^{K}$ – up to discretization with mesh Δt – not with respect to the filtration generated by $(X_{t_k})_{k=0}^{K}$.
- The recursive step to $k = 0$ is reformulated here, to avoid a singular regression.
- Normalization does not seem to matter in practice.
- Replacing the linear regression by a non-linear one possible.

$$V_0 = \inf_{M \in \mathcal{M}_0^2} \mathbb{E}\left[\sup_{t\in[0,T]} (Y_t - M_t)\right]$$

Theorem

Assume that $(\mathcal{F}_t)_{t\in[0,T]}$ *is generated by an* e*-dimensional Brownian motion* W*. Then for any* $M \in \mathcal{M}_0^2$ *there is a sequence* $f_n^1, \dots, f_n^e \in L^\lambda_{\mathit{sig}}(\Lambda_T)$ *s.t.*

$$\int_0^\cdot f_n(\widehat{\mathbb{X}}|_{[0,t]})^\top \mathrm{d}W_t := \sum_{i=1}^e \int_0^\cdot f_n^i(\widehat{\mathbb{X}}|_{[0,t]})\,\mathrm{d}W_t^i \xrightarrow[L^2]{n\to\infty} M.$$

Consequently,

$$V_0 = \inf_{\ell^1,\dots,\ell^e \in \mathcal{W}_{1+d}} \mathbb{E}\left[\sup_{t\in[0,T]}\left(Y_t - \sum_{i=1}^e \int_0^t \left\langle \ell^i, \widehat{\mathbb{X}}_{0,s}^{<\infty}\right\rangle \mathrm{d}W_s^i\right)\right].$$

- Finite number of exercise dates $\mathcal{T} = \{ t_0, \dots, t_K \}$ (Bermudan option)
- Signature $\widetilde{\mathbb{X}}^{\leq N}_{0,t}$ and stochastic integrals $\int_0^t \left\langle \ell^i, \widetilde{\mathbb{X}}^{<\infty}_{0,\lfloor s \rfloor} \right\rangle \mathrm{d}W^i_s$ computed with mesh size Δt.
- Truncation of the signature at level N
- Finite sample size M

Theorem

$$\overline{V}^{\mathcal{T}}_0 \xrightarrow[\Delta t \to 0]{K,N\to\infty} V^{\mathcal{T}}_0 \text{ a.s.}, \qquad \overline{V}^{\mathcal{T}}_0 := \inf_{\ell^1,\dots,\ell^e \in \mathcal{W}^{\leq N}_{1+d}} \frac{1}{M} \sum_{m=1}^{M} \max_{0 \leq k \leq K} \left(Y_{t_k} - \sum_{i=1}^{e} \int_0^{t_k} \left\langle \ell^i, \widetilde{\mathbb{X}}^{<\infty}_{0,\lfloor s \rfloor} \right\rangle \mathrm{d}W^i_s \right)$$

- Finite number of exercise dates $\mathcal{T} = \{ t_0, \dots, t_K \}$ (Bermudan option)
- Signature $\widetilde{\mathbb{X}}^{\le N}_{0,t}$ and stochastic integrals $\int_0^t \left\langle \ell^i, \widetilde{\mathbb{X}}^{<\infty}_{0,\lfloor s \rfloor} \right\rangle \mathrm{d}W^i_s$ computed with mesh size Δt.
- Truncation of the signature at level N
- Finite sample size M

Theorem

$$\overline{V}_0^{\mathcal{T}} \xrightarrow[\Delta t \to 0]{K,N\to\infty} V_0^{\mathcal{T}} \text{ a.s.}, \quad \overline{V}_0^{\mathcal{T}} := \inf_{\ell^1,\dots,\ell^e \in \mathcal{W}^{\le N}_{1+d}} \frac{1}{M} \sum_{m=1}^{M} \max_{0\le k\le K} \left(Y_{t_k} - \sum_{i=1}^{e} \int_0^{t_k} \left\langle \ell^i, \widetilde{\mathbb{X}}^{<\infty}_{0,\lfloor s \rfloor} \right\rangle \mathrm{d}W^i_s \right)$$

- The minimization problem in $\overline{V}_0^{\mathcal{T}}$ is convex, and, in fact, can be formulated as a linear program in dimension $M + \dim \mathcal{W}^{\le N}_{1+d}$.

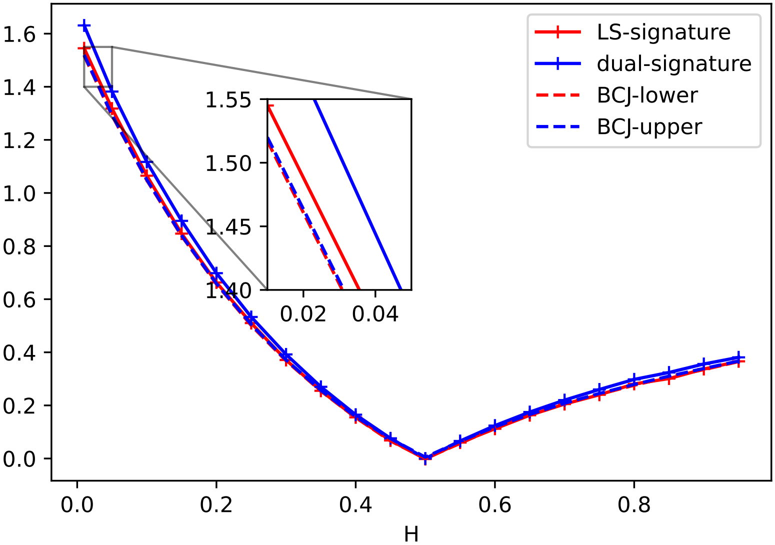

Figure: Approximation based on $J = 500$ time steps, $K = 100$ exercise dates, signature truncation at level $N = 6$. Comparison of the Longstaff–Schwartz and dual martingale methods with results of [Becker-Cheredito-Jentzen '19] based on $J = 100$.

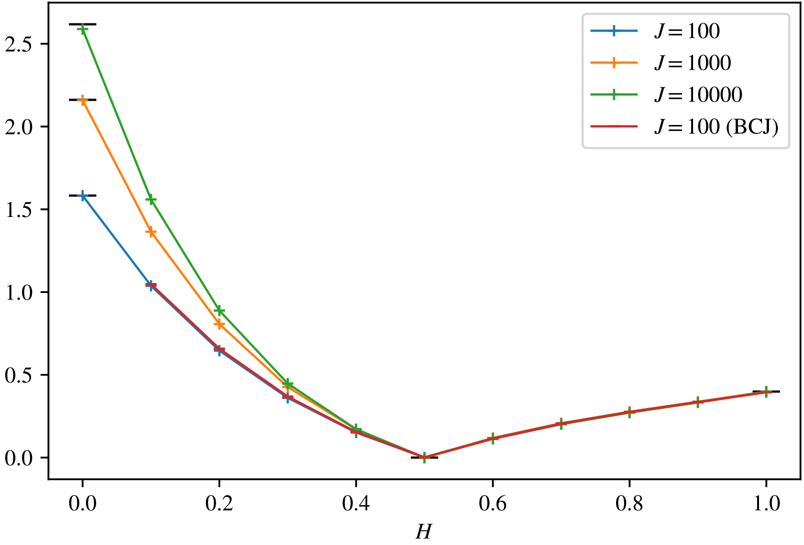

Figure: Approximation based on non-linear parameterization of stopping times in terms of neural networks in the signature, [B., Hager, Riedel, Schoenmakers '23]. Discretization J time steps, log-signature truncated at $N = 3$ ($\dim \mathfrak{g}^{\leq N} = 5$), NN with 2 hidden layers.

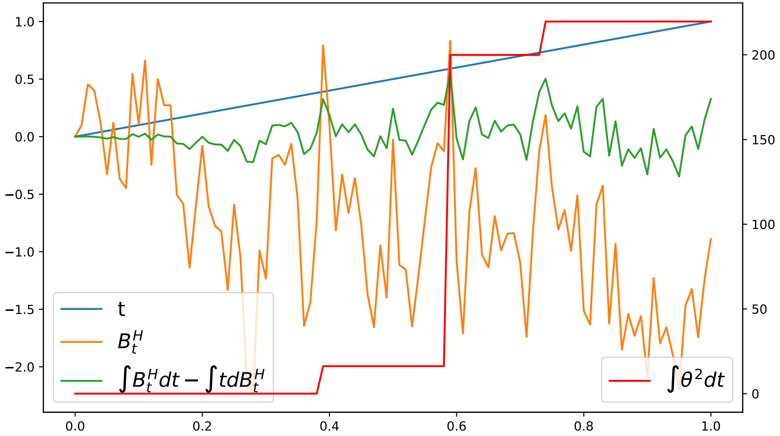

Figure: Approximate randomized stopping rule in [B., Hager, Riedel, Schoenmakers ’23] and selected log-signature entries for one trajectory of a fractional Brownian motion with $H = 0.1$.

K	Lower-bound	Upper-bound	*B. et al., ’20*	*Goudenege et al., ’20*
70	1.92 (± 0.006)	1.99 (± 0.012)	1.88	1.88
80	3.27 (± 0.008)	3.37 (± 0.010)	3.22	3.25
90	5.37 (± 0.011)	5.49 (± 0.012)	5.30	5.34
100	8.57 (± 0.013)	8.77 (± 0.014)	8.50	8.53
110	13.29 (± 0.015)	13.59 (± 0.012)	13.23	13.28
120	20.24 (± 0.013)	20.66 (± 0.010)	20	20.20

Table: Put option prices for the rough Bergomi model with $J = 600$ time steps, $H = 0.07$, truncation at level $N = 3$ for Longstaff–Schwartz (adding polynomials of price and v of degree up to 3) and $N = 4$ for the dual upper bound, $K = 12$ exercise dates.

Thank you for your attention!

- C. Bayer, L. Pelizzari, J. Schoenmakers. *Primal and dual optimal stopping with signatures*. arXiv 2023.
- C. Bayer, P. Hager, S. Riedel, J. Schoenmakers. *Optimal stopping with signatures*. Annals of Appl. Prob., 2023.
- C. Bayer, R. Tempone, and S. Wolfers. *Pricing American options by exercise rate optimization*. Quant. Fin., 2020.
- S. Becker, P. Cheridito, and A. Jentzen. *Deep optimal stopping*. JMLR, 2019.
- C. Cuchiero, P. Schmocker, J. Teichmann. *Global universal approximation of functional input maps on weighted spaces*. arXiv, 2023.
- L. Goudenege, A. Molent, A. Zanette. *Machine learning for pricing American options in high-dimensional Markovian and non-Markovian models*. Quant. Fin., 2020.
- J. Kalsi, T. Lyons, and I. Perez Arribas. *Optimal execution with rough path signatures*. SIAM J. Financial Math., 2020.