Weierstraß-Institut
für Angewandte Analysis und Stochastik
Leibniz-Institut im Forschungsverbund Berlin e. V.

Preprint ISSN 2198-5855

Primal and dual optimal stopping with signatures

Christian Bayer[1], Luca Pelizzari[1,2], John G. M. Schoenmakers[1]

submitted: December 14, 2023

[1] Weierstrass Institute
Mohrenstr. 39
10117 Berlin
Germany
E-Mail: christian.bayer@wias-berlin.de
luca.pelizzari@wias-berlin.de
john.schoenmakers@wias-berlin.de

[2] Institut für Mathematik
Technische Universität Berlin
Fakultät II
Str. des 17. Juni 136
10587 Berlin
Germany
E-Mail: pelizzari@math.tu-berlin.de

No. 3068
Berlin 2023

2020 *Mathematics Subject Classification.* 60L10, 60L20, 91G20, 91G60.

Key words and phrases. Signature, optimal stopping, rough paths, Monte Carlo, rough volatility.

All authors gratefully acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689). The authors would like to thank C. Cuchiero and S. Breneis for valuable remarks and helpful discussions about the global approximation in Section 2.

Edited by
Weierstraß-Institut für Angewandte Analysis und Stochastik (WIAS)
Leibniz-Institut im Forschungsverbund Berlin e. V.
Mohrenstraße 39
10117 Berlin
Germany

Fax: +49 30 20372-303
E-Mail: preprint@wias-berlin.de
World Wide Web: http://www.wias-berlin.de/

Primal and dual optimal stopping with signatures

Christian Bayer, Luca Pelizzari, John G. M. Schoenmakers

Abstract

We propose two signature-based methods to solve the optimal stopping problem – that is, to price American options – in non-Markovian frameworks. Both methods rely on a global approximation result for L^p-functionals on rough path-spaces, using linear functionals of robust, rough path signatures. In the primal formulation, we present a non-Markovian generalization of the famous Longstaff-Schwartz algorithm, using linear functionals of the signature as regression basis. For the dual formulation, we parametrize the space of square-integrable martingales using linear functionals of the signature, and apply a sample average approximation. We prove convergence for both methods and present first numerical examples in non-Markovian and non-semimartingale regimes.

1 Introduction

Stochastic processes with memory play a more and more important role in the modelling of financial markets. In the modelling of equity markets, *rough stochastic volatility models* are now part of the standard toolbox, see, e.g., [23, 3]. In the same area, *path-dependent stochastic volatility models* [26] are a very powerful alternative for capturing memory-effects. Processes with memory are also an essential tool for modelling the micro-structure of financial markets, driven by the market practice of splitting large orders in many medium size ones, as well as by the reaction of algorithmic traders to such orders. Seen from outside, this materializes as self-excitation of the order flow, and, consequently, *Hawkes processes* are a fundamental tool for modelling order flows, see, e.g., [11]. Beyond finance, processes with memory play an important role in the modelling of many natural phenomena (e.g., earthquakes, see [31]) or social phenomena.

In this paper we study optimal stopping problems in non-Markovian frameworks, that is the underlying price is possibly a stochastic process with memory. For concreteness' sake, let us introduce two processes determining the optimal stopping problem: an underlying *state-process* X, together with its natural filtration $\mathcal{F}^X$, and a *reward-process* Z, which is $(\mathcal{F}_t^X)-$adapted – think about $X = (S, v)$ for a stock price process S driven by a stochastic variance process v and $Z_t = \phi(t, S_t)$. The optimal stopping problem then consists of solving the following optimization problem

$$y_0 = \sup_{\tau \in \mathcal{S}_0} E[Z_\tau], \tag{1}$$

where $\mathcal{S}_0$ denotes the set of $(\mathcal{F}_t^X)-$stopping times on $[0, T]$, for some $T > 0$. We merely assume α-Hölder continuity for X in our framework, see Section 3.1 below, in particular allowing non-Markovian and non-semimartingale state-processes X.

The lack of Markov property leads to severe theoretical and computational challenges in the context of optimal control problems, and thus in particular in the optimal stopping problem (1). Indeed,

DOI 10.20347/WIAS.PREPRINT.3068 Berlin 2023

the primary analytical and numerical framework for stochastic optimal control problems arguably is the associated Hamilton–Jacobi–Bellman (HJB) PDE, in the context of optimal stopping so-called *free-boundary problems*, see [32]. When the state process is not a Markov process, such PDEs do not exist a priori. BSDE methods may not be similarly restricted, in principle, but most numerical approximation methods crucially rely on the Markov property, as well.

It should be noted that, at least intuitively, all processes with memory can be turned into Markov processes by adding the history to the current state – but see, e.g., [12] for a more sophisticated approach in the case of fractional Brownian motion. Hence, theoretical and numerical methods from the Markovian world are, in principle, available, but at the cost of having to work in infinite-dimensional (often very carefully drafted, see. e.g., [17]) state spaces. On the other hand, *Markovian approximations*, i.e., finite-dimensional Markov processes closely mimicking the process with memory, can sometimes be a very efficient surrogate model, especially when high accuracy is achievable with low-dimensional Markovian approximations, see, e.g., [2].

Inspired by many successful uses in machine learning (for time-series data), [27] introduced a model-free method for numerically solving a stochastic optimal execution problem. The method is based on the *path signature*, see, e.g., [22], and is applicable in non-Markovian settings. This approach was extended to optimal stopping problems in [4], where stopping times were parameterized as first hitting times of affine hyperplanes in the signature-space. A rigorous mathematical analysis of that method was performed and numerical examples verifying its efficiency were provided.

The signature $\mathbf{X}^{<\infty}$ of a path $X : [0,T] \to \mathbb{R}^d$, is given (at least formally) as the infinite collection of *iterated integrals*, that is for $0 \leq t \leq s \leq T$

$$\mathbf{X}^{<\infty}_{s,t} = \left\{ \int_s^t \int_s^{t_k} \cdots \int_s^{t_2} dX^{i_1}_{t_1} \cdots dX^{i_k}_{t_k} : i_1, \dots, i_k \in \{1, \dots, d\}, k \geq 0 \right\}.$$

The signature characterizes the history of the corresponding trajectory, and, hence, provides a systematic way of “lifting” a process with memory to a Markov process by adding the past to the state. Relying only on minimal regularity assumptions, the corresponding encoding is efficient, and has nice algebraic properties. In many ways, (linear functionals of) the path signature behaves like an analogue of polynomials on path-space, and can be seen as a canonical choice of basis functions on path-space. For example, a *Stone-Weierstrass* type of result shows that, restricted to compacts, continuous functionals on path-spaces can be approximated by linear functionals of the signature, that is by linear combinations of iterated integrals, see for instance [27, Lemma 3.4].

As a first contribution, in Section 2 we provide an abstract approximation result on α-Hölder rough path spaces, by linear functionals of the *robust* signature, with respect to the L^p-norm, see Theorem 2.6 below. As a direct consequence, and under very mild assumptions, we can show that for any $(\mathcal{F}^X_t)-$progressive process $(\xi_t)_{t\in[0,T]}$, we can find a sequence $(l_n)_{n\in\mathbb{N}}$ of linear functionals on the state-space of the signature, such that

$$E\left[\int_0^T (\xi_t - \langle \mathbf{X}^{<\infty}_{0,t}, l_n\rangle)^2 dt\right] \xrightarrow{n\to\infty} 0, \tag{2}$$

see Corollary 2.7 below for the details. This result is in marked contrast to the standard universal approximation result for signatures as usually formulated, which only provides uniform convergence on compact subsets of the path space.

Returning to the optimal stopping problem (1), in Section 3 we generalize two standard techniques from the Markovian case to the non-Markovian case using signatures, namely the *Longstaff–Schwartz algorithm* [29] and Rogers' *dual martingale method* [33]. Denoting by Y the *Snell envelope* to the optimal stopping problem, see below for more details, the Longstaff–Schwartz algorithm is based on the *dynamic programming principle*, that is

$$Y_t = \max\left(Z_t,\ E[Y_{t+\Delta t}|\mathcal{F}_t^X]\right).$$

If X is a Markov process, then the conditional expectation $E[Y_{t+\Delta t}|\mathcal{F}_t^X] = E[Y_{t+\Delta t}|X_t]$, which can be efficiently computed using regression (*least-squares Monte Carlo*). In the non-Markovian case, an application of the global approximation result Theorem 2.6, i.e. the convergence in (2), shows that (under minimal assumptions) a Longstaff–Schwartz algorithm converges when the conditional expectation is approximated by linear functionals of the signature, that is

$$t \mapsto E[Y_{t+\Delta t}|\mathcal{F}_t^X] \approx \langle \mathbf{X}_{0,t}^{<\infty}, l\rangle,$$

see Proposition 3.2.

Regarding the dual method, we rely on Roger's characterization that

$$y_0 = \inf_{M\in\mathcal{M}_0^2} E\left[\sup_{t\in[0,T]} (Z_t - M_t)\right],$$

where the $\inf$ is taken over all square-integrable martingales M starting at 0. If the underlying filtration is Brownian, such martingales can be written as stochastic integrals w.r.t. a Brownian motion W, that is $M_t = \int_0^t \xi_s dW_s$ for some $(\mathcal{F}_t^X)-$progressive process ξ. The approximation result Theorem 2.6, i.e. the convergence in (2), suggests to approximate the integrand by linear functionals of the signature

$$t \mapsto \xi_t \approx \langle \mathbf{X}_{0,t}^{<\infty}, l\rangle,$$

and we prove convergence after taking the infimum over all linear functionals l, that is

$$y_0 = \inf_l E\left[\sup_{t\in[0,T]}\left(Z_t - \int_0^t \langle \mathbf{X}_{0,s}^{<\infty}, l\rangle dW_s\right)\right], \tag{3}$$

see Proposition 3.6. For numerically solving the dual problem (3) we carry out a *Sample Average Approximation* (SAA) with respect to the coefficients of the linear functional of the signature. For a Markovian environment, a related SAA procedure was earlier proposed in [18] and recently refined in [8] and [7] using a suitable randomization. An important feature of the SAA method is that it relies on nonnested Monte Carlo simulation and thus is very fast in comparison to the classical nested Monte Carlo method by Andersen & Broadie [1].

For both the Longstaff–Schwartz and the dual signature methods, we also prove convergence of the finite sample approximations when the number of samples grows to infinity, see Proposition 3.3 and Proposition 3.8. It is worth to notice that, after independent resimulations, the Longstaff–Schwartz algorithm yields lower-biased, whereas the dual method gives upper-biased values to the optimal stopping problem (1), and thus applying both methods produces confidence intervals for the true value of y_0.

Finally, in Section 4 we provide first numerical examples based on the primal and dual signature-based approaches, in two non-Markovian frameworks: First, in Section 4.1, we study the task of optimally stopping *fractional Brownian motion* for a wide range of Hurst-parameter $H \in (0,1)$, representing the canonical choice of a state-process leaving the Markov regime. The same problem was already studied in [6], and later in [4], and we compare our lower, resp. upper bounds with the results therein. Secondly, in Section 4.2 we consider the problem of computing American options prices in the *rough Bergomi model* [3], and we compare our price intervals with [5], resp. [25], where lower-bounds were computed in the same model.

1.1 Notation

For $d, N \in \mathbb{N}$ we define the so-called *extended tensor-algebra*, and the N-step truncation thereof by

$$T((\mathbb{R}^d)) = \prod_{k\geq 0} (\mathbb{R}^d)^{\otimes k}, \quad T^{\leq N}(\mathbb{R}^d) = \prod_{k=0}^{N} (\mathbb{R}^d)^{\otimes k},$$

where we use the convention $(\mathbb{R}^d)^{\otimes 0} = \mathbb{R}$. For more details, including natural operations such as sum $+$ and product $\star$ on these spaces, see for instance [22, Section 7.2.1]. For any word $w = i_1 \cdots i_n$ for some $n \in \mathbb{N}$ with $i_1, \dots, i_n \in \{1, \dots, d\}$, we define the *degree* of w as the length of the word, that is $\deg(w) = n$, and denote by $\emptyset$ the empty word with $\deg(\emptyset) = 0$. Moreover, for $\mathbf{a} \in T((\mathbb{R}^d))$, we denote by $\langle \mathbf{a}, w \rangle$ the element of $\mathbf{a}^{(n)} \in (\mathbb{R}^d)^{\otimes n}$ corresponding to the basis element $e_{i_1} \otimes \cdots \otimes e_{i_n}$. Denoting by $\mathcal{W}^d$ the linear span of words, the pairing above can be extended linearly $\langle \cdot, \cdot \rangle : T((\mathbb{R}^d)) \times \mathcal{W}^d \to \mathbb{R}$. For an element $l \in \mathcal{W}^d$, that is $l = \lambda_1 w_1 + \cdots + \lambda_n w_n$ for some words $w_1, \dots, w_n$ and scalars $\lambda_1, \dots, \lambda_n \in \mathbb{R}$, we define the degree of l by $\deg(l) := \max_{1\leq i\leq n} \deg(w_i)$, and for $K \in \mathbb{N}$ we denote by $\mathcal{W}^d_{\leq K} \subset \mathcal{W}^d$ the subset of elements l with $\deg(l) \leq K$. For two words w and v we denote by $\sqcup\!\sqcup$ the *shuffle-product*

$$w \sqcup\!\sqcup \emptyset = \emptyset \sqcup\!\sqcup w = w, \quad wi \sqcup\!\sqcup vj := (w \sqcup\!\sqcup vj)\, i + (wi \sqcup\!\sqcup v)\, j, \quad i, j \in \{1, \dots, d\}, \tag{4}$$

which bi-linearly extends to the span of words $\mathcal{W}^d$. We further define the *free nilpotent Lie-group* over $\mathbb{R}^d$ by

$$G((\mathbb{R}^d)) = \left\{ \mathbf{a} \in T((\mathbb{R}^d)) \setminus \{\mathbf{0}\} : \langle \mathbf{a}, w \rangle \langle \mathbf{a}, v \rangle = \langle \mathbf{a}, w \sqcup\!\sqcup v \rangle, \forall w, v \in \mathcal{W}^d \right\},$$

see [22, Chapter 7.5] for details.

For $\alpha \in (0,1)$ we denote by $C^\alpha([0,T], \mathbb{R}^d)$ the space of α-Hölder continuous paths X, that is $X : [0,T] \to \mathbb{R}^d$ such that

$$\|X\|_{\alpha;[0,T]} = \sup_{0\leq s<t\leq T} \frac{\|X_t - X_s\|}{|t-s|^\alpha} < \infty,$$

where $\|\cdot\|$ denotes the Euclidean norm on $\mathbb{R}^d$. Denote by $\Delta^2_{[0,T]}$ the simplex $\Delta^2_{[0,T]} := \{(s,t) \in [0,T]^2 : 0 \leq s \leq t \leq T\}$. For any two-parameter function on the truncated tensor-algebra

$$\Delta^2_{[0,T]} \ni (s,t) \mapsto \mathbf{X}_{s,t} = \left(1, \mathbf{X}^{(1)}_{s,t}, \dots, \mathbf{X}^{(N)}_{s,t}\right) \in T^{\leq N}(\mathbb{R}^d),$$

we denote by $|\!|\!|\cdot|\!|\!|_{(\alpha,N)}$ the norm given by

$$|\!|\!|\mathbf{X}|\!|\!|_{(\alpha,N)} := \max_{1\leq k\leq N} \left(\sup_{0\leq s<t\leq T} \frac{\|\mathbf{X}^{(k)}_{s,t}\|}{|t-s|^{k\alpha}} \right)^{1/k} \tag{5}$$

We denote by $\mathscr{C}_g^\alpha([0,T],\mathbb{R}^d)$ the space of geometric $\alpha-$Hölder rough paths $\mathbf{X}$ on $\mathbb{R}^d$, which is the $|||\cdot|||_{(\alpha,N)}-$closure of N-step signatures of Lipschitz continuous paths $X : [0,T] \to \mathbb{R}^d$ for $N = \lfloor 1/\alpha \rfloor$. More precisely, for every $\mathbf{X} \in \mathscr{C}_g^\alpha$ there exists a sequence $(X^n)_{n\in\mathbb{N}} \subset \mathrm{Lip}([0,T],\mathbb{R}^d)$ such that $|||\mathbf{X}^n - \mathbf{X}|||_{(\alpha,N)} \xrightarrow{n\to\infty} 0$, where $\mathbf{X}^n$ is the N-step signature of X^n, that is

$$\mathbf{X}_{s,t}^n := \left(\int_{s<t_1<\cdots<t_k<t} \otimes dX_{t_1}^n \cdots \otimes dX_{t_k}^n : 0 \le k \le N \right) \in G^{\le N}(\mathbb{R}^d),$$

where the integrals are defined in a Riemann-Stieljes sense. For any $\mathbf{X} \in \mathscr{C}_g^\alpha$ we denote by $\mathbf{X}^{<\infty}$ the *rough path signature*, which is the unique (up to tree-like equivalence $\sim_t$[1]) path from Lyon's extension theorem [30, Theorem 3.7], that is

$$\Delta_{[s,t]}^2 : [0,T] \ni (s,t) \mapsto \mathbf{X}_{s,t}^{<\infty} = (1, \mathbf{X}_{s,t}^{(1)}, \dots, \mathbf{X}_{s,t}^{(N)}, \mathbf{X}_{s,t}^{(N+1)}, \dots) \in G(\mathbb{R}^d), \tag{6}$$

such that

$$\|\mathbf{X}^{(k)}\|_{k\alpha} < \infty \quad \forall k \ge 0, \quad \mathbf{X}_{s,t}^{<\infty} = \mathbf{X}_{s,u}^{<\infty} \star \mathbf{X}_{u,t}^{<\infty} \quad s \le u \le t,$$

where the latter is called *Chens* relation. Finally, by considering time-augmented paths $\widehat{X}_t = (t, X_t)$, and their geometric rough path lifts $\widehat{\mathbf{X}}$, the signature maps becomes unique due to the strictly monoton time component. We denote by $\widehat{\mathscr{C}}_g^\alpha([0,T],\mathbb{R}^{d+1})$ the space of geometric α-Hölder rough path lifts of $\widehat{X}_t = (t, X_t)$, where $X \in C^\alpha([0,T],\mathbb{R}^d)$.

2 Global approximation with rough path signatures

In this section we present the theoretical foundation of this paper, which consists of a global approximation result based on *robust* rough path signatures.

2.1 The space of stopped rough paths

For $\alpha \in (0,1)$, we consider an α-Hölder continuous path $X : [0,T] \to \mathbb{R}^d$ with $X_0 = x_0 \in \mathbb{R}^d$, and denote by $\mathbf{X}$ the geometric rough path lift of (t, X_t), that is $\mathbf{X} \in \widehat{\mathscr{C}}_g^\alpha([0,T],\mathbb{R}^{d+1})$.

Definition 2.1 *For any $\alpha \in (0,1)$ and $T > 0$, the space of stopped $\widehat{\mathscr{C}}_g^\alpha-$paths is defined by the disjoint union*

$$\Lambda_T^\alpha := \bigcup_{t\in[0,T]} \widehat{\mathscr{C}}_g^\alpha([0,t],\mathbb{R}^{d+1}).$$

Moreover, we equip the space Λ_T^α with the final topology[2] *induced by the map*

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

The reason to work on this space is the following: If X is a stochastic process, and $\mathbf{X}$ denotes the random geometric lift of (t, X_t), we define $\mathcal{F}_t^{\mathbf{X}} = \sigma(\mathbf{X}_{0,s} : s \le t)$ for $0 \le t \le T$, i.e. the natural filtration generated by $\mathbf{X}$. In Lemma 2.4 below, we show that any $(\mathcal{F}_t^{\mathbf{X}})-$progressive process

[1] $\sim_t$ is an equivalence class on path-spaces, including for example time-changes, see [9] for details.

[2] Recall that for a topological space Y and $f : Y \to X$, the final topology on X, induced by f, consists of all sets $A \subseteq X$ s.t. $f^{-1}(A)$ is open.

DOI 10.20347/WIAS.PREPRINT.3068

$(A_t)_{t\in[0,T]}$ can be expressed as $A_t = f(\mathbf{X}|_{[0,t]})$, where f is a measurable function on Λ_T^α. Thus, progressively measurable processes can be thought of measurable functionals on Λ_T^α, and we will discuss approximation results for the latter below. Similar spaces have already been considered in relation with *functional Itô calculus* in [19, 15], and for p-rough paths in [27], and more recently in relation with optimal stopping in [4].

Remark 2.2 *One can also introduce a metric d_Λ on the space Λ_T^α, defined by*

$$d_\Lambda(\mathbf{x}|_{[0,t]}, \mathbf{y}|_{[0,s]}) := |||\mathbf{x} - \tilde{\mathbf{y}}|||_{(\alpha,N):[0,t]} + |t-s|, \quad s \le t,$$

where $\tilde{\mathbf{y}}$ is the rough path lift uf $u \mapsto (u, y_{u\wedge s})$ for $u \in [0,t]$. It has been proved in [4, Lemma A.1], in the case of p-rough paths, the topology of the metric space (Λ_T^p, d_Λ) coincides with the final topology, and the space of stopped geometric rough paths is Polish. A similar argument can be done for the α-Hölder case, by replacing the p-variation norm by the α-Hölder norm and using the fact that $\widehat{\mathscr{C}}_g^\alpha$ is Polish, see [22, Proposition 8.27].

Remark 2.3 *Let $\mathbf{X}$ be a stochastic process on $\widehat{\mathscr{C}}_g^\alpha([0,T],\mathbb{R}^{d+1})$ for all $\alpha < \gamma \in (0,1)$. It is discussed in [22, Appendix A.1] that $\mathbf{X}$ can be regarded as $\widehat{\mathscr{C}}_g^\alpha([0,T],\mathbb{R}^{d+1})-$valued random variable, and its law $\mu_{\mathbf{X}}$ is a Borel measure on the Borel σ-algebra $\mathcal{B}^\alpha$ with respect to $|||\cdot|||_{(\alpha,N)}$. Moreover, define the product measure $d\mu := dt \otimes d\mu_{\mathbf{X}}$. For the surjection $\phi : [0,T] \times \widehat{\mathscr{C}}_g^\alpha([0,T],\mathbb{R}^{d+1}) \to \Lambda_T^\alpha$ defined above, we can define the pushforward measure $\widehat{\mu}$ on Λ_T^α, in symbols $\widehat{\mu} := \phi\#\mu$, which is given by*

$$\widehat{\mu}(A) := \mu(\phi^{-1}(A)) \text{ for all } A \in \mathcal{B}(\Lambda_T^\alpha). \tag{7}$$

Consider the space $\mathbb{H}^2$ of $(\mathcal{F}_t^{\mathbf{X}})-$progressive processes A, such that

$$\|A\|_{\mathbb{H}^2}^2 := E\left[\int_0^T A_s^2 ds\right] < \infty. \tag{8}$$

The following result justifies the consideration of the space Λ_T^α.

Lemma 2.4 *For any $A \in \mathbb{H}^2$ and $\alpha < \gamma \in (0,1)$, there exists a measurable function $f : (\Lambda_T^\alpha, \mathcal{B}(\Lambda_T^\alpha)) \longrightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$, such that $A_t = f(\mathbf{X}|_{[0,t]})$ almost everywhere.*

Proof Consider the space of elementary, $(\mathcal{F}_t^{\mathbf{X}})-$progressive processes, that is processes of the form

$$A_t^n(\omega) := \xi_0^n(\omega)1_{\{0\}}(t) + \sum_{j=1}^{m_n-1} 1_{(t_j^n, t_{j+1}^n]}(t)\xi_{t_j}^n(\omega), \tag{9}$$

where $0 \le t_0^n < \cdots < t_{m_n}^n \le T$, and $\xi_{t_j}^n$ is a $(\mathcal{F}_{t_j}^{\mathbf{X}})-$measurable, square integrable random variable. A standard result for the construction of stochastic integrals, shows that this space is dense in $\mathbb{H}^2$, this can be found in [28, Lemma 3.2.4] for instance. Thus, we can find A^n of the form (9), such that $A^n \longrightarrow A$ for almost every (t,ω). Notice that for all $\xi_{t_j}^n$ there exists a Borel measurable function $F_j^n : \widehat{\mathscr{C}}_g^\alpha([0,t_j],\mathbb{R}^{d+1}) \longrightarrow \mathbb{R}$ such that $\xi_{t_j}^n(\omega) = F_j^n(\mathbf{X}|_{[0,t_j]}(\omega))$. Then the functions

$$[0,T] \times \widehat{\mathscr{C}}_g^\alpha \ni (t,\mathbf{x}) \mapsto 1_{(t_j^n, t_{j+1}^n]}(t) F_j^n(\mathbf{x}|_{[0,t_j]}) \tag{10}$$

are $\left(\mathcal{B}([0,T]) \otimes \mathcal{F}_T^{\mathbf{X}}\right)-$measurable, and therefore also the function

$$F^n(t,\mathbf{x}) := F_0^n(\mathbf{x}_0)1_{\{0\}}(t) + \sum_{j=1}^{m_n-1} 1_{(t_j^n, t_{j+1}^n]}(t) F_j^n(\mathbf{x}|_{[0,t_j]}).$$

Finally, define the jointly measurable function $F(t,\mathbf{x}) := \limsup_{n\to\infty} F^n(t,\mathbf{x})$, and notice that for almost every (t,ω), we have

$$F(t,\mathbf{X}(\omega)) = \limsup_{n\to\infty}\left(F_0^n(\mathbf{X}_0(\omega))1_{\{0\}}(t) + \sum_{j=1}^{m_n-1} 1_{(t_j^n,t_{j+1}^n]}(t)F_j^n(\mathbf{X}|_{[0,t_j]}(\omega))\right)$$
$$= \limsup_{n\to\infty}\left(\xi_0^n(\omega)1_{\{0\}}(t) + \sum_{j=1}^{m_n-1} 1_{(t_j^n,t_{j+1}^n]}(t)\xi_{t_j}^n(\omega)\right) = A_t(\omega).$$

Next, for any element $\mathbf{x}|_{[0,t]} \in \Lambda_T^\alpha$, we let $(t,\tilde{\mathbf{x}}) \in [0,T]\times\widehat{\mathscr{C}}_g^\alpha([0,T],\mathbb{R}^{d+1})$, where $\tilde{\mathbf{x}}$ is the geometric rough path lift of $[0,T] \ni u \mapsto (u, x_{u\wedge t})$. The map $\Xi : \Lambda_T^\alpha \longrightarrow [0,T]\times\widehat{\mathscr{C}}_g^\alpha$ with $\Xi(\mathbf{x}|_{[0,t]}) := (t,\tilde{\mathbf{x}})$ is continuous and thus especially measurable. Define the composition $f := F\circ\Xi$, which is a measurable map $f : \Lambda_T^\alpha \longrightarrow \mathbb{R}$, such that

$$f(\mathbf{X}|_{[0,t]}(\omega)) = A_t(\omega),$$

for almost every (t,ω), which is exactly what was claimed. □

2.2 A Stone–Weierstrass result for robust signatures

The goal of this section is to present a Stone-Weierstrass type of result for continuous functionals $f : \Lambda_T^\alpha \longrightarrow \mathbb{R}$, which will be the key ingredient for the main result in Section 2.3. To this end, consider the set of linear functionals of the signature

$$L_{\mathrm{sig}}(\Lambda_T^\alpha) = \left\{\Lambda_T^\alpha \ni \mathbf{X}|_{[0,t]} \mapsto \langle \mathbf{X}_{0,t}^{<\infty}, l\rangle : l \in \mathcal{W}^{d+1}\right\} \subseteq C(\Lambda_T^\alpha,\mathbb{R}),$$

where we recall that $\mathcal{W}^{d+1}$ denotes the linear span of words, see Section 1.1. A similar set was considered in [27, Definition 3.3] with respect to p-rough paths, and the authors prove that *restricted* to a compact set K on the space of time-augemented rough paths, the set L_{sig} is dense in $C(K,\mathbb{R})$. In words, restricted to compacts, continuous functionals on the path-space Λ_T^α can be approximated by linear functionals of the signature. However, since such path-spaces are not even locally compact, it is desirable to drop the need of a compact set K.

An elegant way to circumvent the requirement of a compact set, is to consider so-called *robust signatures*, introduced in [13]. Loosely speaking, the authors construct a so-called *tensor-normalization* λ, see [13, Proposition 14 and Example 4], on the state-space of the signature $T((\mathbb{R}^{d+1}))$, which is a continuous and injective map

$$\lambda : T((\mathbb{R}^{d+1})) \longrightarrow \{\mathbf{a}\in T((\mathbb{R}^{d+1})) : \|\mathbf{a}\| \le R\}, \quad R>0,$$

and they call $\lambda(\mathbf{X}^{<\infty})$ the robust signature. This motivates to define the set

$$L_{\mathrm{sig}}^\lambda(\Lambda_T^\alpha) = \left\{\Lambda_T^\alpha \ni \mathbf{X}|_{[0,t]} \mapsto \langle \lambda(\mathbf{X}_{0,t}^{<\infty}), l\rangle : l \in \mathcal{W}^{d+1}\right\} \subseteq C_b(\Lambda_T^\alpha,\mathbb{R}).$$

A general version of the Stone-Weierstrass result given in [24], leads to the following result, which was stated already in [13, Theorem 26], and we present the proof here for completeness.

Lemma 2.5 *Let $\alpha \in (0,1)$. Then the set $L^{\lambda}_{\mathrm{sig}}(\Lambda^{\alpha}_{T})$ is dense in $C_b(\Lambda^{\alpha}_{T}, \mathbb{R})$ with respect to the strict topology. More precisely, for any $f \in C_b(\Lambda^{\alpha}_{T}, \mathbb{R})$ we can find a sequence $(f_n)_{n\in\mathbb{N}} \subseteq L^{\lambda}_{\mathrm{sig}}(\Lambda^{\alpha}_{T})$, such that*

$$\|f - f_n\|_{\infty,\psi} := \sup_{\mathbf{x}\in\Lambda^{\alpha}_{T}} |\psi(\mathbf{x})(f(\mathbf{x}) - f_n(\mathbf{x}))| \xrightarrow{n\to\infty} 0, \quad \forall \psi \in B_0(\Lambda^{\alpha}_{T}),$$

where $B_0(\Lambda^{\alpha}_{T})$ denotes the set of functions $\psi : \Lambda^{\alpha}_{T} \to \mathbb{R}$, such that for all $\epsilon > 0$ there exists a compact set $K \subseteq \Lambda^{\alpha}_{T}$ with $\sup_{\mathbf{x}\in\Lambda^{\alpha}_{T}\setminus K} |\psi(\mathbf{x})| < \epsilon$.

Proof This result is a consequence of the general Stone-Weierstrass result proved in [24], see also [13, Theorem 9]. From the latter, we only need to check that $L^{\lambda}_{\mathsf{sig}} \subseteq C_b(\Lambda^{\alpha}_{T}, \mathbb{R})$ is a subalgebra, such that

1 $L^{\lambda}_{\mathsf{sig}}$ separates points, that is $\forall x \neq y$ there exists $f \in L^{\lambda}_{\mathsf{sig}}$ such that $f(x) \neq f(y)$.

2 $L^{\lambda}_{\mathsf{sig}}$ contains non-vanishing functions, that is $\forall x$ there exists $f \in L^{\lambda}_{\mathsf{sig}}$ such that $f(x) \neq 0$.

To see that $L^{\lambda}_{\mathrm{sig}} \subseteq C_b(\Lambda^{\alpha}_{T}, \mathbb{R})$ is a subalgebra, we fix $\phi_1, \phi_2 \in L^{\lambda}_{\mathrm{sig}}$. By definition, there exist words $l_1, l_2 \in \mathcal{W}^{d+1}$ such that $\phi_i(\mathbf{X}|_{[0,s]}) = \langle \lambda(\mathbf{X}^{<\infty}_{0,s}), l_i\rangle$. We clearly have

$$\phi(\mathbf{X}|_{[0,s]}) := \phi_1(\mathbf{X}|_{[0,s]}) + \phi_2(\mathbf{X}|_{[0,s]}) = \langle \lambda(\mathbf{X}^{<\infty}_{0,s}), l_1 + l_2\rangle \in L^{\lambda}_{\mathrm{sig}}.$$

Now assume l_1, l_2 are words $l_1 = w$ and $l_2 = v$. By definition of the tensor-normalization [13, Definition 12], for some positive function $\Psi : T((\mathbb{R}^{d+1})) \to]0, +\infty[$, we have

$$\begin{aligned}\phi_1(\mathbf{X}|_{[0,s]}) \cdot \phi_2(\mathbf{X}|_{[0,s]}) &= \langle \lambda(\mathbf{X}^{<\infty}_{0,s}), w\rangle \langle \lambda(\mathbf{X}^{<\infty}_{0,s}), v\rangle \\ &= \Psi(\mathbf{X}^{<\infty}_{0,s})^{|w|+|v|} \langle \mathbf{X}^{<\infty}_{0,s}, w\rangle \langle \mathbf{X}^{<\infty}_{0,s}, v\rangle \\ &= \Psi(\mathbf{X}^{<\infty}_{0,s})^{|w|+|v|} \langle \mathbf{X}^{<\infty}_{0,s}, w \sqcup\!\sqcup v\rangle,\end{aligned}$$

where we used that $\mathbf{X}^{<\infty} \in G(\mathbb{R}^d)$ for the last equality. But by definition of the shuffle-product (4), it follows that $w \sqcup\!\sqcup v = \sum_j u_j$, where u_j are words with $|u_j| = |w| + |v|$, and hence

$$\phi_1(\mathbf{X}|_{[0,s]}) \cdot \phi_2(\mathbf{X}|_{[0,s]}) = \sum_j \Psi(\mathbf{X}^{<\infty}_{0,s})^{|w|+|v|} \langle \mathbf{X}^{<\infty}_{0,s}, u_j\rangle = \langle \lambda(\mathbf{X}^{<\infty}_{0,s}), w \sqcup\!\sqcup v\rangle \in L^{\lambda}_{\mathrm{sig}}.$$

The same reasoning can be extended to linear combination of words l_1, l_2, and thus the set $L^{\lambda}_{\mathrm{sig}}$ is indeed a subalgebra in $C_b(\Lambda^{\alpha}_{T}, \mathbb{R})$. Now let $\mathbf{X}, \mathbf{Y} \in \Lambda^{\alpha}_{T}$, such that $\mathbf{X} \neq \mathbf{Y}$. As remarked in Section 1.1, since we are working with rough path lifts of time-augmented paths (t, X_t), the signature map is injective. Moreover, by definition [13, Definition 12], the map λ is also injective, therefore $\lambda(\mathbf{X}^{<\infty}) \neq \lambda(\mathbf{Y}^{<\infty})$ and thus $L^{\lambda}_{\mathrm{sig}}$ separates points. Finally, since $1 = \langle \lambda(\mathbf{X}^{<\infty}), \emptyset\rangle \in L^{\lambda}_{\mathrm{sig}}$, the claim follows. □

2.3 Approximation with robust signatures

We are now ready to state and proof the main result of this section. For a fixed $\gamma \in (0,1)$, we consider the Borel space $(\Lambda^{\alpha}_{T}, \mathcal{B}(\Lambda^{\alpha}_{T}))$ for $\alpha < \gamma$, as described in Section 2.1. Suppose μ is a measure on $(\Lambda^{\alpha}_{T}, \mathcal{B}(\Lambda^{\alpha}_{T}))$ such that

$$\mu\left(\Lambda^{\alpha}_{T}\right) < \infty \text{ and } \mu\left(\Lambda^{\alpha}_{T} \setminus \Lambda^{\beta}_{T}\right) = 0, \quad \forall \beta \in (\alpha, \gamma). \tag{11}$$

The following theorem shows that under assumption (11), we can approximate any functional in $L^p(\Lambda^{\alpha}_{T}, \mu)$ by linear functionals of the robust signature with respect to the L^p-norm.

Theorem 2.6 *Let $\alpha < \gamma \in (0,1)$ and consider the measure space $(\Lambda_T^\alpha, \mathcal{B}(\Lambda_T^\alpha), \mu)$ such that* (11) *holds true. Then for all $f \in L^p(\Lambda_T^\alpha, \mu)$, $1 \le p < \infty$, there exists a sequence $(f_n)_{n\in\mathbb{N}} \subset L^\lambda_{sig}(\Lambda_T^\alpha)$ such that $\|f - f_n\|_{L^p} \xrightarrow{n\to\infty} 0$.*

Before proving this result, let us show the following immediate consequence for random geometric rough paths, which will be of particular importance in Section 3.

Corollary 2.7 *Let $\mathbf{X}$ be a stochastic process on $\widehat{\mathscr{C}}_g^\alpha$ for $\alpha < \gamma$. Then for all $A \in \mathbb{H}^2$, see* (8)*, there exists a sequence $(f_n)_{n\in\mathbb{N}} \subset L^\lambda_{sig}(\Lambda_T^\alpha)$ such that $\|A - f_n(\mathbf{X}|_{[0,\cdot]})\|_{\mathbb{H}^2} \xrightarrow{n\to\infty} 0$.*

Proof Consider the measure $\widehat{\mu}$ on Λ_T^α defined in Remark 2.3, that is the push-forward of the product measure $dt\otimes d\mu_{\mathbf{X}}$ on $[0,T]\times\widehat{\mathscr{C}}_g^\alpha$, which is a finite measure that assigns full measure to the subspaces Λ_T^β for all $\beta \in (\alpha,\gamma)$. Thus (11) holds true, and by Lemma 2.4 we know that there exists a measurable function $f : \Lambda_T^\alpha \to \mathbb{R}$ such that $A_t = f(\mathbf{X}|_{[0,t]})$. Applying a standard change of measure result for the push-forward measure, see for example [10, Theorem 3.6.1], and denoting by ϕ the quotient map given in Definition 2.1, we have

$$\begin{aligned}\|A_t - f_n(\mathbf{X}|_{[0,\cdot]})\|_{\mathbb{H}^2}^2 &= E\left[\int_0^T (f(\mathbf{X}|_{[0,t]}) - f_n(\mathbf{X}|_{[0,t]}))^2 dt\right] \\ &= \int_{\widehat{\mathscr{C}}_g^\alpha}\int_0^T (f\circ\phi - f_n\circ\phi)(t,\mathbf{X})^2 dt d\mu_{\mathbf{X}} \\ &= \int_{\Lambda_T^\alpha} (f-f_n)^2 d\widehat{\mu} = \|f-f_n\|_{L^2}^2 \xrightarrow{n\to\infty} 0,\end{aligned}$$

where the convergence follows from Theorem 2.6 for $p = 2$. □

The proof of Theorem 2.6 will make use of two lemmas. The first one is very elementary, and in the language of probability theory, it states that for every random variable X in $\mathbb{R}_+$, we can find a strictly increasing and integrable function η, that is $E[\eta(X)] < \infty$.

Lemma 2.8 *Let $(E,\mathcal{E},\mu)$ be a finite measure space, and $\xi : (E,\mathcal{E}) \to (\mathbb{R}_+, \mathcal{B}(\mathbb{R}_+))$ a measurable function. Then there exists a strictly increasing function $\eta : \mathbb{R}_+ \to \mathbb{R}_+$, such that $\eta(x) \xrightarrow{x\to\infty} \infty$ and $\int_E (\eta\circ\xi)d\mu < \infty$.*

Proof Let ν be the push-foward of μ under ξ, that is $\nu(A) := \mu(\xi^{-1}(A))$ for all $A \in \mathcal{B}(\mathbb{R}_+)$. Then for any $\epsilon > 0$ we can find $R > 0$ large enough, such that $\nu(]R,+\infty[) \le \epsilon$. In particular, for any strictly decreasing sequence $(a_n)_{n\ge 0}$, such that $a_n \searrow 0$, we can find a strictly increasing sequence $(R_n)_{n\in\mathbb{N}}$ with $R_1 > 0$, such that $\nu(]R_n,\infty[) \le \frac{a_n}{n^2}$. Now we can define a strictly increasing function η as follows: Let $\eta(0) = 0$ and for all $n \in \mathbb{N}$ define $\eta(R_n) = \frac{1}{a_{n-1}}$, and linearly interpolate on the intervals $[R_n, R_{n+1}[$. Then, setting $R_0 = 0$, and using a change of measure [10, Theorem 3.6.1], we have

$$\int_E (\eta\circ\xi)d\mu = \int_0^\infty \eta d\nu \le \sum_{n\ge 0}\frac{1}{a_n}\nu(]R_n,+\infty[) \le \frac{1}{a_0} + \sum_{n\ge 1}\frac{1}{n^2} < \infty.$$

□

The next lemma will be the key ingredient to apply the Stone-Weierstrass result Lemma 2.5 in the main result.

Lemma 2.9 *Let $\alpha < \beta$ and define the function $\bar\psi(\mathbf{x}) := 1_{\Lambda_T^\beta}(\mathbf{x})\left(\frac{1}{1+\eta(|||\tilde{\mathbf{x}}|||_{(\beta,N)})}\right)^{1/p}$, where $\eta : \mathbb{R}_+ \to \mathbb{R}_+$ is strictly increasing such that $\eta(x) \xrightarrow{x\to\infty} \infty$. Then $\bar\psi \in B_0(\Lambda_T^\alpha)$, that is for all $\epsilon > 0$ there exists $K \subseteq \Lambda_T^\alpha$ compact, such that $\sup_{\mathbf{x}\in K^c}\psi(\mathbf{x}) \le \epsilon$.*

Proof Recall that an element $\mathbf{x} \in \Lambda_T^\alpha$ can be written as $\mathbf{x} = \mathbf{x}|_{[0,t]} \in \widehat{\mathscr{C}}_g^\alpha([0,t])$, and $\mathbf{x}|_{[0,t]}$ is the rough path lift of some time-augmented, α-Hölder continuous path $[0,t] \ni u \mapsto (u,\omega_u)$. Moreover, recall that we define $\tilde{\mathbf{x}} \in \widehat{\mathscr{C}}_g^\alpha([0,T])$ to be the geometric rough path lift of $u \mapsto (u, x_{t\wedge u})$. If we can show that for any $R > 0$, the sets

$$B_R = \{\mathbf{x} \in \Lambda_T^\beta : |||\tilde{\mathbf{x}}|||_{(\beta,N)} \le R\} \subseteq \Lambda_T^\alpha$$

are compact, then we are done. Indeed, in this case we have that for any $\epsilon \in (0,1)$, we can choose $\widehat{R} \ge \eta^{-1}(\frac{1-\epsilon^p}{\epsilon^p})$ and then

$$\psi(\mathbf{x}) = \left(\frac{1}{1+\eta(|||\tilde{\mathbf{x}}|||_{(\beta,N)})}\right)^{1/p} \le \frac{1}{\sqrt{1+\widehat{R}}} \le \epsilon, \quad \forall \mathbf{x} \in \Lambda_T^\alpha \setminus B_{\widehat{R}},$$

and therefore $\psi \in B_0(\Lambda_T^\alpha)$. Now to prove compactness, we can notice that by definition of the quotient map ϕ, see Definition 2.1, we have

$$B_R \subseteq \phi\left([0,T] \times \left\{\mathbf{x} \in \widehat{\mathscr{C}}_g^\beta([0,T],\mathbb{R}^{d+1}) : |||\mathbf{x}|||_{(\beta,N)} \le R\right\}\right),$$

since for all $\mathbf{x} = \mathbf{x}|_{[0,t]} \in B_R$ we have $\mathbf{x} = \phi(t,\tilde{\mathbf{x}})$ by construction. Since ϕ is continuous, it is enough the show that $[0,T] \times \{\mathbf{x} \in \widehat{\mathscr{C}}_g^\beta : |||\mathbf{x}|||_{(\beta,N)} \le R\}$ is compact in $[0,T] \times \widehat{\mathscr{C}}_g^\alpha$, which by Tychonoffs theorem is true if the sets $\{\mathbf{x} \in \widehat{\mathscr{C}}_g^\beta : |||\mathbf{x}|||_{(\beta,N)} \le R\}$ are compact in $\widehat{\mathscr{C}}_g^\alpha$. But the latter follows from the general fact that β-Hölder spaces are compactly embedded in α-Hölder spaces for $\alpha < \beta$. This can be proved by applying the Arzelà–Ascoli theorem together with an interpolation argument for the equicontinuous and $|||\cdot|||_{(\beta,N)}$-bounded subsets of $\widehat{\mathscr{C}}_g^\alpha$, which was carried out in [16, Theorem A.3] for example. Thus we can conclude that $B_R \subseteq \Lambda_T^\alpha$ is compact, which finishes the proof. □

Finally, we are ready to proof the main result.

Proof of Theorem 2.6 Fix $\epsilon > 0$. For any $K > 0$, we can define the function $f_K(x) := 1_{\{|f(x)|\le K\}}(x)f(x)$, and notice that we have $\|f - f_K\|_{L^p} \to 0$ as $K \to \infty$ by dominated convergence. Hence we can find a $K_\epsilon > 0$ such that $\|f - f_{K_\epsilon}\|_{L^p} \le \epsilon/3$. Since μ is a finite measure on Λ_T^α, by Lusin's theorem, we can find a closed set $C_\epsilon \subset \Lambda_T^\alpha$, such that f_{K_ϵ} restricted to C_ϵ is continuous, and $\mu(\Lambda_T^\alpha \setminus C_\epsilon) \le \epsilon^p/(6K_\epsilon)^p$. By Tietze's extension theorem, we can find a continuous extension $\widehat{f}_\epsilon \in C_b(\Lambda_T^\alpha, [-K_\epsilon, K_\epsilon])$ of f_{K_ϵ} such that

$$\|f_{K_\epsilon} - \widehat{f}_\epsilon\|_{L^p}^p = \int_{\Lambda_T^\alpha\setminus C_\epsilon} |f_\epsilon - f_{K_\epsilon}|^p d\mu \le (2K_\epsilon)^p \mu(\Lambda_T^\alpha \setminus C_\epsilon) = (\epsilon/3)^p.$$

We are left with approximating $\widehat{f}_\epsilon \in C_b(\Lambda_T^\alpha, \mathbb{R})$ by linear functionals of the robust signature, that is applying Lemma 2.5. From Lemma 2.8 we know, for any $\beta < \gamma$ we know there exists an increasing function $\eta : \mathbb{R}_+ \to \mathbb{R}_+$, such that $\eta(x) \xrightarrow{x\to\infty} \infty$ and $\int_{\Lambda_T^\alpha} \eta(|||\tilde{\mathbf{x}}|||_{(\beta,N)}) d\mu(x) < \infty$, where $\tilde{\mathbf{x}}$ is the extension of the stopped rough path from the interval $[0,t]$ to $[0,T]$, see also Remark 2.2. Fix $\beta \in (\alpha,\gamma)$ and define the function $\psi : \Lambda_T^\beta \to \mathbb{R}_+$ by $\psi(\mathbf{x}|_{[0,t]}) := \left(\frac{1}{1+\eta(|||\tilde{\mathbf{x}}|||_{(\beta,N)})}\right)^{1/p}$. Defining

$\Xi := \int_{\Lambda_T^\beta} \frac{1}{\psi^p} d\mu$, it follows from Lemma 2.8 that $\Xi < \infty$. In Lemma 2.9 we saw that $\bar{\psi}(\mathbf{x}) := 1_{\Lambda_T^\beta}(\mathbf{x})\psi(\mathbf{x})$ belongs to $B_0(\Lambda_T^\alpha)$, that is for all $\delta > 0$ there exists a compact set $K \subseteq \Lambda_T^\alpha$, such that $\sup_{x\in K^c} \bar{\psi}(x) \le \delta$. By Lemma 2.5, we can find $f_\epsilon \in L^\lambda_{\mathrm{sig}}(\Lambda_T^\alpha)$, such that

$$\|\widehat{f}_\epsilon - f_\epsilon\|^p_{\infty,\bar{\psi}} \le \epsilon^p/(3^p\Xi).$$

Using that μ assigns full measure to the subspace $\Lambda_T^\beta \subseteq \Lambda_T^\alpha$, we have

$$\begin{aligned}\|\widehat{f}_\epsilon - f_\epsilon\|^p_{L^p} = \int_{\Lambda_T^\beta} |\widehat{f}_\epsilon - f_\epsilon|^p d\mu &\le \sup_{x\in\Lambda_T^\beta} \Big(\psi(x)(\widehat{f}_\epsilon(x) - f_\epsilon(x))\Big)^p \int_{\Lambda_T^\beta} \frac{1}{\psi^p} d\mu \\ &\le \|\widehat{f}_\epsilon - f_\epsilon\|^p_{\infty,\bar{\psi}}\Xi \le (\epsilon/3)^p.\end{aligned}$$

Finally, we can conclude by the triangle inequality

$$\|f - f_\epsilon\|_{L^p} \le \|f - f_{K_\epsilon}\|_{L^p} + \|f_{K_\epsilon} - \widehat{f}_\epsilon\|_{L^p} + \|\widehat{f}_\epsilon - f_\epsilon\|_{L^p} \le \epsilon.$$

□

3 Optimal stopping with signatures

In this section we exploit the signature approximation theory presented in Section 2.3, in order solve the optimal stopping problem in a general setting.

3.1 Framework and problem formulation

Let $T > 0$ and consider a complete, filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in[0,T]}, P)$ fulfilling the usual conditions, and fix $\gamma \in (0,1)$. For any $(\mathcal{F}_t)-$adapted and α-Hölder continuous stochastic process $(X_t)_{t\in[0,T]}$, $\alpha < \gamma$, taking values in $\mathbb{R}^d$ with $X_0 = x_0$, we consider

- $\mathbf{X} \in \widehat{\mathscr{C}}_g^\alpha$ the geometric α-Hölder rough path lift of (t, X_t) for $\alpha < \gamma$,
- $\mathbf{X}^{<\infty}$ the robust[3] rough path signature introduced in Section 2.3,
- $(Z_t)_{t\in[0,T]}$ is a real-valued, $(\mathcal{F}_t^{\mathbf{X}})-$adapted stochastic process such that $\sup_{t\in[0,T]} |Z_t| \in L^2$.

The optimal stopping problem then reads

$$y_0 = \sup_{\tau\in\mathcal{S}_0} E[Z_\tau], \tag{12}$$

where $\mathcal{S}_0$ denotes the set of $(\mathcal{F}_t^{\mathbf{X}})$-stopping-times on $[0,T]$.

Remark 3.1 *Notice that the framework described above is very general in two ways: First, we only assume α-Hölder continuity for the state process X, including in particular non-Markovian and non-semimartingale regimes, which one for instance encounters in rough volatility models, see Section 4.2. Secondly, considering a projection onto the first coordinate $\mathbf{X} \mapsto (t, X_t)$, for any payoff function $\phi : [0,T]\times\mathbb{R}^d \to \mathbb{R}$ our framework includes the more common form of the optimal stopping problem*

$$y_0 = \sup_{\tau\in\mathcal{S}_0} E[\phi(\tau, X_\tau)].$$

[3]Notice the small abuse of notation here, as the robust signature is given by $\lambda(\mathbf{X}^{<\infty})$ for some tensor-normalization λ. For the rest of this paper, we fix such an λ and write $\mathbf{X}^{<\infty}$ for $\lambda(\mathbf{X}^{<\infty})$.

3.2 Primal optimal stopping with signatures

First we present a method to compute a lower-biased approximation $y_0^L \leq y_0$ to the optimal stopping problem (12). More precisely, we construct a regression-based approach, generalizing the famous algorithm from Longstaff and Schwartz [29], returning a sub-optimal exercise strategy. Let us first quickly describe the main idea of most regression-based approaches.

Replacing the interval $[0,T]$ by a finite grid $\{0 = t_0 < t_1 < \cdots < t_N = T\}$, the discrete optimal stopping problem reads

$$y_0^N = \sup_{\tau \in \mathcal{S}_0^N} E[Z_\tau],$$

where $\mathcal{S}_0^N$ is the set of stopping times taking values in $\{t_0, \dots, t_N\}$, with respect to the discrete filtration $(\mathcal{F}_{t_n}^{\mathbf{X}})_{n=0,\dots,N}$. We define the discrete *Snell-envelope* by

$$Y_{t_n}^N = \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_n^N} E[Z_\tau | \mathcal{F}_{t_n}^{\mathbf{X}}], \quad 0 \leq n \leq N, \tag{13}$$

and one can show that Y^N satisfies the discrete dynamic programming principle (DPP)

$$Y_{t_n}^N = \max\left(Z_{t_n}, E[Y_{t_{n+1}}^N | \mathcal{F}_{t_n}^{\mathbf{X}}]\right), \quad n = 0, \dots, N-1, \tag{14}$$

see for instance [32, Theorem 1.2]. Now the key idea of most regression-based approaches, such as for instance [29], is that assuming X is a Markov process, one can choose a suitable family of basis functions (b^k) and apply least-square regression to approximate

$$E[Y_{t_{n+1}}^N | \mathcal{F}_{t_j}^{\mathbf{X}}] \approx \sum_{k=0}^{D} \alpha_k b_n^k(X_{t_n}), \quad 0 \leq n \leq N-1, \quad \alpha_k \in \mathbb{R}, \quad \forall k \leq D, \tag{15}$$

and then make use of the DPP to recursively approximate $Y_0^N = y_0^N$. Of course, the approximation of the conditional expectations in (15) heavily relies on the Markov-property, and thus one cannot expect such an approximation to converge in non-Markovian settings.

Returning to our framework, we need to replace (15) by a suitable approximation for the conditional expectations

$$E[Y_{t_{n+1}}^N | \mathcal{F}_{t_n}^{\mathbf{X}}] = f_n(\mathbf{X}|_{[0,t_n]}), \quad 0 \leq n \leq N-1. \tag{16}$$

The universality result Theorem 2.6 now suggests to approximate f_n by a sequence of linear functionals of the robust signature, that is

$$f_n(\mathbf{X}|_{[0,t_n]}) \approx \langle \mathbf{X}_{0,t_n}^{<\infty}, l \rangle, \quad l \in \mathcal{W}^{d+1}, \tag{17}$$

where $\mathcal{W}^{d+1}$ is the linear span of words introduced in Section 1.1.

Longstaff-Schwartz with signatures

In this section we present a version of the Longstaff-Schwartz (LS) algorithm [29], using signature-based least-square regression. A convergence analysis for the LS-algorithm was presented in [14], and combining their techniques with the universality of the signature, allows us to recover a convergent algorithm.

The main idea of the LS-algorithm is to re-formulate the DPP (14) for stopping times, taking advantage of the fact that optimal stopping times can be expressed in terms of the Snell-envelope. More precisely, it is proved in [32, Theorem 1.2] that the stopping times $\tau_n := \min\{t_m \geq t_n : Y^N_{t_m} = Z_{t_m}, m = n, \dots, N\}$ are optimal in (13), and hence one recursively defines

$$\begin{cases} \tau_N = t_N \\ \tau_n = t_n 1_{\left\{Z_{t_n} \geq E\left[Z_{\tau_{n+1}} | \mathcal{F}^{\mathbf{X}}_{t_n}\right]\right\}} + \tau_{n+1} 1_{\left\{Z_{t_n} < E\left[Z_{\tau_{n+1}} | \mathcal{F}^{\mathbf{X}}_{t_n}\right]\right\}}, \quad n = 0, \dots, N-1. \end{cases} \tag{18}$$

Now for any truncation level $K \in \mathbb{N}$ for the signature, we approximate the conditional expectations in (18) by solving the following minimization problem

$$l^* := l^{*,n,K} = \underset{l \in \mathcal{W}^{d+1}_{\leq K}}{\operatorname{argmin}} \left\| Z_{\tau_{n+1}} - \langle \mathbf{X}^{\leq K}_{0,t_n}, l \rangle \right\|_{L^2}, \quad n = 0, \dots, N-1. \tag{19}$$

Setting $\psi^{n,K}(\mathbf{x}) = \langle \mathbf{x}^{\leq K}, l^{*,n,K} \rangle \in L^{\lambda}_{\text{sig}}$, we can define the following approximating sequence of stopping times

$$\begin{cases} \tau^K_N = t_N \\ \tau^K_n = t_n 1_{\left\{Z_{t_n} \geq \psi^{n,K}(\mathbf{X}|_{[0,t_n]})\right\}} + \tau^K_{n+1} 1_{\left\{Z_{t_n} < \psi^{n,K}(\mathbf{X}|_{[0,t_n]})\right\}}, \quad n = 0, \dots, N-1. \end{cases} \tag{20}$$

The following result shows convergence as the depth of the signature goes to infinity, and the proof is discussed in Appendix A.1.

Proposition 3.2 *For all* $n = 0, \dots, N$ *we have*

$$E[Z_{\tau^K_n} | \mathcal{F}^{\mathbf{X}}_{t_n}] \xrightarrow{K \to \infty} E[Z_{\tau_n} | \mathcal{F}^{\mathbf{X}}_{t_n}] \text{ in } L^2.$$

In particular, we have $y^{K,N}_0 = \max(Z_{t_0}, E[Z_{\tau^K_1}]) \xrightarrow{K \to \infty} y^N_0$.

Let us now describe how to numerically solve (19) using Monte-Carlo simulations. Besides the truncation of the signature at some level K, we introduce two further approximations steps: First, we replace the signature $\mathbf{X}^{<\infty}$ by some discretized version[4] $\mathbf{X}^{<\infty}(J)$ on some fine grid $s_0 = 0 < s_1 < \dots < s_J = T$, such that $E[\langle \mathbf{X}^{<\infty}_{0,t}(J), v \rangle] \xrightarrow{J \to \infty} E[\langle \mathbf{X}^{<\infty}_{0,t}, v \rangle]$ for all words v and $t \in [0, T]$. Secondly, for $i = 1, \dots, M$ i.i.d sample paths of Z and the discretized and truncated signature $\mathbf{X}^{\leq K} = \mathbf{X}^{\leq K}(J)$, assuming that $\tau^{K,J}_{n+1}$ is known, we estimate l^* by solving (19) via linear least-square regression. This yields an estimator $l^* = l^{*,n,J,K,M}$. Defining $\psi^{n,J,K,M}(\mathbf{x}) = \langle \mathbf{x}^{\leq K}, l^{*,n,J,K,M} \rangle$ leads to a recursive algorithm for stopping times, for $i = 1, \dots, M$

$$\begin{cases} \tau^{K,J,(i)}_N = t_N \\ \tau^{K,J,(i)}_n = t_n 1_{\left\{Z^{(i)}_{t_n} \geq \psi^{n,J,K,M}(\mathbf{X}^{(i)}|_{[0,t_n]})\right\}} + \tau^{K,J,(i)}_{n+1} 1_{\left\{Z^{(i)}_{t_n} < \psi^{n,J,K,M}(\mathbf{X}^{(i)}|_{[0,t_n]})\right\}}. \end{cases} \tag{21}$$

Then the following law of large number type of result holds true, which almost directly follows from [14, Theorem 3.2], see Appendix A.1.

Proposition 3.3 *For all* $n = 0, \dots, N$ *we have*

$$\frac{1}{M} \sum_{i=1}^{M} Z^{(i)}_{\tau^{K,J,(i)}_n} \xrightarrow{M,J \to \infty} E[Z_{\tau^K_n}] \text{ a.s.}$$

[4] For instance piecewise linear approximation of the iterated integrals.

In particular, we have

$$y_0^{K,N,J,M} := \max\left(Z_{t_0}, \frac{1}{M}\sum_{i=1}^{M} Z^{(i)}_{\tau_1^{K,J,(i)}}\right) \xrightarrow{M,J\to\infty} y_0^{K,N} \text{ a.s.},$$

and thus especially $|y_0^{K,N,J,M} - y_0^N| \xrightarrow{M,J,K\to\infty} 0.$

Remark 3.4 *The recursion of stopping times* (21)*, resp. the resulting linear functionals of the signature* $\psi^{n,K,M}$*, provide a stopping policy for each sample path of* Z*. By resimulating* $\tilde{M}$ *i.i.d samples of* Z *and the signature* $\mathbf{X}^{<\infty}$*, we can notice that the resulting estimator* $y_0^{K,N,J,\tilde{M}}$ *is lower-biased, that is* $y_0^{K,N,J,\tilde{M}} \leq y_0^N$*, since the latter is defined by taking the supremum over all possible stopping policies.*

3.3 Dual optimal stopping with signatures

In this section, we approximate solutions to the optimal stopping problem in its dual formulation, leading to upper bounds $y_0^U \geq y_0$ for (12). The dual representation goes back to [33], where the author shows that under the assumption $\sup_{0\leq t\leq T}|Z_t| \in L^2$, the optimal stopping problem (12) is equivalent to

$$y_0 = \inf_{M\in\mathcal{M}_0^2} E\left[\sup_{t\leq T}(Z_t - M_t)\right], \tag{22}$$

where $\mathcal{M}_0^2$ denotes the space of $(\mathcal{F}_t^{\mathbf{X}})$-martingales in L^2, starting from 0. Assuming that $\mathcal{F}^{\mathbf{X}}$ is generated by a Brownian motion W, we can prove the following equivalent formulation of (22).

Theorem 3.5 *Assume that* $\mathcal{F}^{\mathbf{X}}$ *is generated by a* m*-dimensional Brownian motion* W*. Then for all* $M \in \mathcal{M}_0^2$*, there exist sequences* $l^i = (l_n^i)_{n\in\mathbb{N}} \subset \mathcal{W}^{d+1}$ *for* $i = 1,\dots,m$*, such that*

$$\int_0^\cdot \langle \mathbf{X}_{0,s}^{<\infty}, l_n\rangle^\top dW_s := \sum_{i=1}^m \int_0^\cdot \langle \mathbf{X}_{0,s}^{<\infty}, l_n^i\rangle dW_s^i \xrightarrow{n\to\infty} M. \text{ ucp.}$$

In particular, the minimization problem (22) *can be equivalently formulated as*

$$\begin{aligned} y_0 &= \inf_{l\in(\mathcal{W}^{d+1})^m} E\left[\sup_{t\leq T}\left(Z_t - \int_0^t \langle \mathbf{X}_{0,s}^{<\infty}, l\rangle^\top dW_s\right)\right] \\ &= \inf_{l^1,\dots,l^m\in\mathcal{W}^{d+1}} E\left[\sup_{t\leq T}\left(Z_t - \sum_{i=1}^m\int_0^t \langle \mathbf{X}_{0,s}^{<\infty}, l^i\rangle dW_s^i\right)\right]. \end{aligned} \tag{23}$$

Proof By the Martingale Representation Theorem, see for instance [28, Theorem 4.5], any $(\mathcal{F}_t^{\mathbf{X}})-$martingale can be represented as

$$M_t = \int_0^t \alpha_s^\top dW_s = \sum_{i=1}^m \int_0^t \alpha_s^i dW_s^i,$$

where $(\alpha_s)_{s\in[0,T]}$ is $(\mathcal{F}_t^{\mathbf{X}})-$adapted, measurable and square integrable. Moreover, since $M \in \mathcal{M}_0^2$, it follows that $E[M_T^2] = E\left[\int_0^T |\alpha_t|^2 dt\right] < \infty$. From [28, Proposition 1.1.12], we know that any adapted and measurable process has a progressively measurable modification, which we again denote by α. From Corollary 2.7, we know that there exist sequences $l^i = (l_n^i)_{n\in\mathbb{N}} \subset \mathcal{W}^{d+1}$ for

$i = 1, \dots, m$, such that for $\alpha_t^{i,n} := \langle \mathbf{X}_{0,t}^{<\infty}, l_n^i \rangle$, we have $\|\alpha^{i,n} - \alpha^i\|_{\mathbb{H}^2} \longrightarrow 0$ as $n \to \infty$. Using Doobs inequality, we in particular have

$$E\left[\left(\sup_{t\le T}\int_0^t (\alpha_s^n - \alpha_s)^\top dW_s\right)^2\right] \lesssim \sum_{i=1}^m E\left[\int_0^T |\alpha_s^{i,n} - \alpha_s^i|^2 dt\right] = \sum_{i=1}^m \|\alpha^{i,n} - \alpha^i\|_{\mathbb{H}^2}^2 \longrightarrow 0.$$

But this readily implies the first claim, that is

$$\int_0^{\cdot} \langle \mathbf{X}_{0,s}^{<\infty}, l_n\rangle^\top dW_s \longrightarrow \int_0^{\cdot} \alpha_s dW_s = M. \text{ ucp.}$$

In order to show (23), since $\int_0^{\cdot} \langle \mathbf{X}_{0,s}^{<\infty}, l_n\rangle^\top dW_s$ are clearly $(\mathcal{F}_t^{\mathbf{X}})$-martingales, we can notice that

$$\inf_{l\in(\mathcal{W}^{d+1})^m} E\left[\sup_{t\le T}\left(Z_t - \int_0^t \langle \mathbf{X}_{0,s}^{<\infty}, l\rangle^\top dW_s\right)\right] \ge \inf_{M\in\mathcal{M}_0^2} E\left[\sup_{t\le T}(Z_t - M_t)\right] = y_0.$$

On the other hand, for any fixed martingale M, we know there exist sequences $l^i = (l_n^i)_{n\in\mathbb{N}} \subset \mathcal{W}^{d+1}$ such that

$$\sup_{t\le T}\left(Z_t - \int_0^t \langle \mathbf{X}_{0,s}^{<\infty}, l_n\rangle^\top dW_s\right) \xrightarrow{n\to\infty} \sup_{t\le T}(Z_t - M_t) \text{ in } L^2.$$

Therefore

$$E\left[\sup_{t\le T}(Z_t - M_t)\right] = \lim_{n\to\infty} E\left[\sup_{t\le T}\left(Z_t - \int_0^t \langle \mathbf{X}_{0,s}^{<\infty}, l_n\rangle^\top dW_s\right)\right] \ge \inf_{l\in(\mathcal{W}^{d+1})^m} E\left[\sup_{t\le T}\left(Z_t - \int_0^t \langle \mathbf{X}_{0,s}^{<\infty}, l\rangle^\top dW_s\right)\right].$$

Taking the infimum over all $M \in \mathcal{M}_0^2$ yields the claim. □

Next, similar to the primal case, we translate the minimization problem (23) into a finite-dimensional optimization problem, by discretizing the interval $[0, T]$ and truncating the signature to some level K. More precisely, for $0 = t_0 < \cdots < t_N = T$ and some $K \in \mathbb{N}$, we reduce the minimization problem (23) to

$$y_0^{K,N} = \inf_{l\in\left(\mathcal{W}_{\le K}^{d+1}\right)^m} E\left[\max_{0\le n\le N}\left(Z_{t_n} - M_{t_n}^l\right)\right], \tag{24}$$

where for any $l = (l^1, \dots, l^m) \in \left(\mathcal{W}_{\le K}^{d+1}\right)^m$ we define

$$M_t^l = \int_0^t \langle \mathbf{X}_{0,s}^{\le K}, l\rangle^\top dW_s = \sum_{i=1}^m \int_0^t \langle \mathbf{X}_{0,s}^{\le K}, l^i\rangle dW_s^i.$$

The discrete version of the dual formulation (22) is given by

$$y_0^N = \inf_{M\in\mathcal{M}_0^{2,N}} E\left[\max_{0\le n\le N}(Z_{t_n} - M_{t_n})\right], \tag{25}$$

where $\mathcal{M}_0^{2,N}$ denotes the space of discrete $(\mathcal{F}_{t_n}^{\mathbf{X}})_{n=0}^N$−martingales. The following result shows that the minimization problem (24) has a solution and the optimal value converges to y_0^N as the level of the signature goes to infinity, the proof can be found in Appendix A.2.

Proposition 3.6 *There exists a minimizer $l^\star$ to* (24) *and*

$$|y_0^N - y_0^{K,N}| \longrightarrow 0 \quad \text{as } K \to \infty.$$

Remark 3.7 *In a financial context, Proposition 3.2 and 3.6 tell us that $y_0^{K,N}$ converges to the Bermudda option price as $K \to \infty$. Moreover, we can use the triangle inequality to find*

$$|y_0 - y_0^{K,N}| \leq |y_0 - y_0^N| + |y_0^N - y_0^{K,N}|, \tag{26}$$

and hence the finite-dimensional approximations converge to y_0 as $K, N \to \infty$, whenever the Bermuddan price converges to the American price. For our numerical examples, we will always approximate y_0^N for some fixed N, and therefore we do not further investigate in the latter convergence here.

Sample average approximation (SAA)

We now present a method to approximate the value $y_0^{K,N}$ in (24), using Monte-Carlo simulations. This procedure is called *sample average approximation* (SAA) and we refer to [34, Chapter 6] for a general and extensive study of this method. Similar to the dual case, we introduce two further approximation steps: First, let $0 = s_0 < \cdots < s_J = T$ be a finer discretization of $[0, T]$ and denote by $M^{l,J}$, resp. $\mathbf{X}^{<\infty}(J)$, an approximation of the stochastic integral $\int \langle \mathbf{X}_{0,s}^{<\infty}, l \rangle dW_s$, resp. the signature $\mathbf{X}^{<\infty}$, using an Euler-scheme. Secondly, we consider $i = 1, \dots, M$ i.i.d. sample paths $Z^{(i)}, M^{(i),l,J}$, and replace the expectation in (24) by a sample average, leading to the following empirical minimization problem

$$y_0^{K,N,J,M} = \inf_{l \in \left(\mathcal{W}_{\leq K}^{d+1}\right)^m} \frac{1}{M} \sum_{i=1}^{M} \max_{0 \leq n \leq N} \left(Z_{t_n}^{(i)} - M_{t_n}^{(i),l,J} \right). \tag{27}$$

The following result can be deduced from [34, Chapter 6 Theorem 4], combined with Proposition 3.6, we refer to Appendix A.2 for the details.

Proposition 3.8 *For M large enough there exists a minimizer $\beta^\star$ to* (27) *and*

$$|y_0^{K,N,J,M} - y_0^N| \longrightarrow 0 \text{ as } K, J, M \longrightarrow \infty.$$

Remark 3.9 *Let us quickly describe how we will solve* (27) *numerically: Consider the number $D := \sum_{k=0}^{K} (d+1)^k$, which corresponds to the number of entries of the K-step signature. Notice that we can write any word $l \in \mathcal{W}_{\leq K}^{d+1}$ as $l = \lambda_1 w_1 + \cdots + \lambda_D w_D$, where $w_1, \dots, w_D$ are all possible words of length at most K. Since $\langle \mathbf{X}_{0,t}^{\leq K}, l \rangle = \sum_{r=1}^{D} \lambda_r \langle \mathbf{X}_{0,t}^{\leq K}, w_r \rangle$, the minimization* (27) *has equivalent formulation*

$$y_0^{K,N,J,M} = \inf_{\lambda \in (\mathbb{R}^D)^m} \frac{1}{M} \sum_{i=1}^{M} \max_{0 \leq n \leq N} \left(Z_{t_n}^{(i)} - \sum_{r=1}^{D} \lambda_r M_{t_n}^{(i),w_r,J} \right).$$

As described in [18], the latter minimization problem is equivalent to the following linear program

$$\min_{x \in \mathbb{R}^{M+D}} \frac{1}{M} \sum_{j=1}^{M} x_j, \quad \text{subject to Ax} \geq b, \tag{28}$$

where $A \in \mathbb{R}^{M(N+1) \times (M+D)}$ with $A = [A^1, \dots, A^M]^T$ and $Ax \geq b$ represents the constraints

$$x_i \geq Z_{t_n}^{(i)} - \sum_{r=1}^{D} M_{t_n}^{(i),w_r,J}, \quad \begin{array}{l} i = 1, \dots, M \\ n = 0, \dots, N \end{array}.$$

Remark 3.10 *A solution $l^\star$ to* (27) *yields the* $(\mathcal{F}^{\mathbf{X}}_{t_n})-$*martingale* $M^{l^\star}$*, and by resimulating* $\tilde{M}$ *i.i.d samples of* Z*, the Brownian motion* W *and the signature* $\mathbf{X}^{<\infty}$*, we can notice that the resulting estimator* $y_0^{K,N,J,\tilde{M}}$ *is upper-biased, that is* $y_0^{K,N,J,\tilde{M}} \geq y_0^N$*, since the latter is defined by taking the infimum over all* $(\mathcal{F}^{\mathbf{X}}_{t_n})-$*martingales.*

4 Numerical examples

In this section we study two non-Markovian optimal stopping problems and test our methods to approximate lower and upper bounds for the optimal stopping value.

4.1 Optimal stopping of fractional Brownian motion

We start with the task of optimally stopping a fractional Brownian motion (fBm), which represents the canonical choice of a framework leaving the Markov and semimartingale regimes. Recall that the fBm with Hurst parameter $H \in (0,1)$ is the unique, continuous Gaussian process $(X_t^H)_{t\in[0,T]}$, with

$$E[X_t^H] = 0, \quad \forall t \geq 0,$$
$$E[X_s^H X_t^H] = \frac{1}{2}\left(|s|^{2H} + |t|^{2H} - |t-s|^{2H}\right), \quad \forall s,t \geq 0.$$

see for instace [21, Chapter 9] for more details. We wish to approximate the value

$$y_0^H = \sup_{\tau\in\mathcal{S}_0} E[X_\tau^H], \quad H \in (0,1), \tag{29}$$

from below and above. This example has already been studied in [6, Section 4.3] as well as in [4, Section 8.1], and we compare the results below.

Since X^H is one-dimensional and α-Hölder continuous for any $\alpha < H$, its (scalar) rough path lift is given by

$$\left(1, X_{s,t}^H, \frac{1}{2}(X_{s,t}^H)^2, \ldots, \frac{1}{N!}(X_{s,t}^H)^N\right) \in \mathscr{C}_g^\alpha([0,T],\mathbb{R}),$$

and we can extend it to a geometric rough path lift $\mathbf{X}^H \in \widehat{\mathscr{C}}_g^\alpha([0,T],\mathbb{R}^2)$ of the time-augmentation (t, X_t^H), as for instance described in [4, Example 2.4]. To numerically solve (29), we replace the interval $[0,T]$ by some grid $0 = t_0 < t_1 < \cdots < t_N = T$. Below we compare our results with [6, Section 4.3], where the authors chose $N = 100$. Before doing so, an important remark about the difference of our problem formulation is in order.

Remark 4.1 *In [6] the authors lift* X^H *to a* 100*-dimensional Markov process of the form* $\widehat{X}_{t_k} = (X_{t_k}^H, \ldots, X_{t_1}^H, 0, \ldots, 0) \in \mathbb{R}^{100}$*, and they consider the corresponding discrete (!) filtration* $\widehat{\mathcal{F}}_k = \sigma(X_{t_k}^H, \ldots, X_{t_1}^H), k = 0, \ldots, 100$*. Notice that this differs from our setting, as we consider the bigger filtration* $\mathcal{F}_k = \sigma(X_s^H : s \leq t_k)$*, see Section 3.2 and 3.3, that contains the whole past of* X^H*, not only the information at the past exercise-dates. Thus, in general* y_0^H *dominates the lower-bounds from [6], simply because our filtration contains more stopping-times. Similarly, the (very sharp) upper-bounds in [6] were obtained using a nested Monte-Carlo approach, which constructs* $(\widehat{\mathcal{F}}_k)$*-martingales that are not martingales in our filtration, and thus their upper-bounds are not necessarily upper-bounds for* (29)*.*

H	$J = 100$	$J = 500$	Becker et al. [6]
0.01	[1.518,1.645]	[1.545,1.631]	[1.517,1.52]
0.05	[1.293,1.396]	[1.318,1.382]	[1.292,1.294]
0.1	[1.045,1.129]	[1.065,1.117]	[1.048,1.05]
0.15	[0.83,0.901]	[0.847,0.895]	[0.838,0.84]
0.2	[0.654,0.706]	[0.663,0.698]	[0.657,0.659]
0.25	[0.507,0.538]	[0.510,0.533]	[0.501,0.505]
0.3	[0.363,0.396]	[0.371,0.392]	[0.368,0.371]
0.35	[0.248,0.272]	[0.255,0.270]	[0.254,0.257]
0.4	[0.153,0.168]	[0.155,0.165]	[0.154,0.158]
0.45	[0.069,0.077]	[0.068,0.076]	[0.066,0.075]
0.5	[-0.001,0]	[-0.002,0]	[0,0.005]
0.55	[0.061,0.071]	[0.060,0.066]	[0.057,0.065]
0.6	[0.112,0.133]	[0.112,0.124]	[0.115,0.119]
0.65	[0.163,0.187]	[0.163,0.175]	[0.163,0.166]
0.7	[0.203,0.234]	[0.205,0.220]	[0.206,0.208]
0.75	[0.242,0.273]	[0.240,0.260]	[0,242,0.245]
0.8	[0.275,0.306]	[0.281,0.298]	[0.276,0.279]
0.85	[0.306,0.335]	[0.301,0.324]	[0.307,0.31]
0.9	[0.331,0.357]	[0.337,0.356]	[0.335,0.339]
0.95	[0.367,0.381]	[0.366,0.381]	[0.365,0.367]

Table 1: Intervals for optimal stopping of fBm $H \mapsto y_0^H$ with $N = 100$ exercise-dates, and discretization $J = 100$ (left column), $J = 500$ (middle column), and intervals from [6] (right column). Overall Monte-Carlo error is below 0.003.

In Table 1 we present intervals for the optimal stopping values y_0^H for $H \in \{0.01, 0.05, \ldots, 0.95\}$, where the lower-bounds, resp. the upper-bounds, were approximated using Longstaff-Schwartz with signatures, resp. the SAA approach described in Section 3. We truncate the signature at level $K = 6$, and apply the primal approach using $M = 10^6$ samples for both the regression and the resimulation, and for the dual approach we choose $M = 15000$ to solve the linear programm from Remark 3.9, and resimulate with $M = 10^5$ samples. In the first column, we choose the time-discretization for the signature equal to the number of exercise-dates, by $J = N = 100$. While the lower-bounds are very close, our upper-bounds exceed the ones from [6]. This observation matches with the comments made in Remark 4.1, as we consider the filtration $\widehat{\mathcal{F}}$ in this case for the lower-bounds, but our upper-bounds are by construction upper-bounds for the continuous problem with filtration $\mathcal{F}$, and the continuous martingale is approximated only at the exercise-dates. By increasing the discretization to $J = 500$, and thereby adding information to the filtration in-between exercise-dates, for small H (≤ 0.2), one can see that the lower-bounds exceed the intervals from [6], showing that even for $N = 100$ points in $[0,1]$, the information in-between exercise-dates is relevant for optimally stopping the fBm.

4.2 American options in rough volatility models

The second example we present is the problem of pricing American options in rough volatility models. More precisely, we consider the one-dimensional asset-price model

$$X_0 = x_0, \quad dX_t = rX_t dt + X_t v_t \left(\rho dW_r + \sqrt{1-\rho^2} dB_t\right), \quad 0 < t \leq T, \tag{30}$$

where W and B are two independent Brownian motions, the volatility $(v_t)_{t\in[0,T]}$ is an $(\mathcal{F}_t^W)-$adapted, continuous process, $\rho \in [-1,1]$ and $r > 0$ the interest rate. Now for any payoff function $\phi : [0,T] \times \mathbb{R} \to \mathbb{R}$, we want to approximate the optimal stopping problem

$$y_0 = \sup_{\tau \in \mathcal{S}_0} E[e^{-r\tau}\phi(\tau, X_\tau)], \tag{31}$$

where $\mathcal{S}_0$ is the set of $(\mathcal{F}_t) := (\mathcal{F}_t^W \vee \mathcal{F}_t^B)-$stopping times on $[0,T]$. It is worth to note that our method does not depend on the specification of v, and as soon as we can sample from (X, v), we can apply it to approximate values of American options.

In the following numerical experiments, we will consider two different signatures. First, since X is one-dimensional, we already saw in Section 4.1 how to construct the rough path signature $\mathbf{X}^{<\infty}$ for the time-augemented path (t, X_t). Since $\mathcal{F}^X = \mathcal{F}^{\mathbf{X}}$, and $Z_t = e^{-rt}\phi(t, X_t) = f(\mathbf{X}_{0,t})$, we are in the setting described in the last section, see also Remark 3.1. Secondly, we do not change the latter framework when adding elements to the path, that is lifting (t, X_t, Q_t) for some path Q, as long as its rough path signature is well-defined. Indeed, numerical experiments suggest to add the payoff process $Q = Z$ in the dual-problem. Since (t, X_t, Z_t) is a semimartingale, the signature $\mathbf{Z}^{<\infty}$ is given as the sequence of iterated Stratonovich integrals as explained in [20]. Finally, we add polynomials of the states (X_t, v_t) to the family of basis-functions in both the primal and dual approach, which of course does not change the convergence. To summarize, for the least squares regression (19), resp. for the SAA minimization problem in (24), we use basis functions of the form

$$\begin{aligned} &\textbf{(P)} \quad \left\{L_i(X_t, v_t), \langle \mathbf{X}_{0,t}^{<\infty}, l\rangle : i = 1, \ldots, m_p, l \in \mathcal{W}_{\leq K_p}^{d+1}\right\}, \quad m_p, K_p \in \mathbb{N}, \\ &\textbf{(D)} \quad \left\{L_i(X_t, v_t), \langle \mathbf{Z}_{0,t}^{<\infty}, l\rangle : i = 1, \ldots, m_d, l \in \mathcal{W}_{\leq K_d}^{d+1}\right\}, \quad m_d, K_d \in \mathbb{N}, \end{aligned} \tag{32}$$

K	Lower-bound	Upper-bound	*Bayer et. al [5]*	*Goudenege et. al [25]*
70	1.85 (± 0.007)	2.04 (± 0.022)	1.88	1.88
80	3.18 (± 0.009)	3.44 (± 0.028)	3.22	3.25
90	5.25 (± 0.011)	5.68(± 0.060)	5.30	5.34
100	8.44 (± 0.013)	9.13 (± 0.014)	8.50	8.53
110	13.18 (± 0.014)	14.18 (± 0.022)	13.23	13.28
120	20.22 (± 0.012)	21.40 (± 0.018)	20	20.20

Table 2: Put option prices with $J = 48$ and $H = 0.07$.

where $(L_k)_{k\geq 0}$ are Laguerre polynomials.

In the following we focus on the rough Bergomi model [3], that is we specify the volatility as

$$v_t = \xi_0 \mathcal{E}\left(\eta \int_0^t (t-s)^{H-\frac{1}{2}} dW_s\right),$$

where $\mathcal{E}$ denotes the stochastic exponential, and we will consider the parameters $x_0 = 100, r = 0.05, \eta = 1.9, \rho = -0.9, \xi_0 = 0.09$. Moreover, we consider put options $\psi(t,x) = (K-x)^+$ for different strikes $K \in \{70, 80, \dots, 120\}$, with maturity $T = 1$ and $N = 12$ exercise-dates. Thus we write (31) as discrete optimal stopping problem

$$y_0^N = \sup_{\tau \in \mathcal{S}_0} E[e^{-r\tau}(K - X_\tau)^+], \tag{33}$$

where S_0^N is described in the beginning of Section 3.2. Moreover, for some finer grid $s_0 = 0 < s_1 < \cdots < s_J = 1, J \in \mathbb{N}$ and fixed signature level K and sample size M, we denote by y_0^{LS} the value $y_0^{K,N,J,M}$ defined in Proposition 3.3, resp. by y_0^{SAA} the value $y_0^{K,N,J,M}$ defined in (27). We compare our results with the lower-bounds obtained in [5] for $H = 0.07$, resp. in [25] with $H = 0.07$ and $H = 0.8$.

In Table 2-3 we compare their lower bounds with our price intervals $[y_0^{\mathsf{LS}}, y_0^{\mathsf{SAA}}]$ together with the Monte-Carlo errors for $H = 0.07$, and different discretizations J. For the lower bound, we fix $K = 3$ for the truncated signature and add polynomials of degree up to 3, that is $m_p = 15$ in **(P)**, and apply the algorithm described in Section 3.2 for $M = 10^6$ samples. For the obtained stopping policies, we resimulate with again $M = 10^6$ samples to obtain true-lower bounds y_0^{LS}. For the upper-bounds, we consider again $K = 4$, and polynomials up to degree 5, that is $m_d = 40$[5], and solve the linear programm in the (SAA) approach described in Section 3.3 for $M = 10^4$ samples, and then resimulate with $M = 10^5$ samples to obtain true upper-bounds y_0^{SAA}. We can observe that similar as in Section 4.1, the discretization plays an important role for small values of H, and the price intervals shrink as we increase the number of discretization points between the exercise-dates. Our lower-bounds exceed the reference values for $J = 600$, and the upper-bounds are $2\% - 3\%$ higher than the lower-bounds. We expect these margins to shrink more when further increasing all the parameters, but we reached the limit of our computational possiblities in Table 3. In Table 4 we consider the same problem for $H = 0.8$ for $J = 600$.

[5]Recall that we consider the lift of (t, X_t, Z_t) for the dual-problem, therefore the number of signature entries is higher compared to the primal case.

K	Lower-bound	Upper-bound	*Bayer et. al [5]*	*Goudenege et. al [25]*
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 3: Put option prices with $J = 600$ and $H = 0.07$.

K	Lower-bound	Upper-bound	*Goudenege et. al [25]*
70	1.83 (± 0.008)	1.90 (± 0.012)	1.84
80	3.08 (± 0.011)	3.19 (± 0.014)	3.10
90	5.07 (± 0.012)	5.17 (± 0.015)	5.08
100	8.15 (± 0.013)	8.27 (± 0.013)	8.19
110	12.97 (± 0.013)	13.09 (± 0.013)	13.00
120	20.21 (± 0.013)	20.51 (± 0.016)	20.28

Table 4: Put option prices $J = 600$ and $H = 0.8$.

A Technical details Section 3

A.1 Proofs in Section 3.2

Proof of Proposition 3.2 The proof is based on the same ideas as the proof in [14, Theorem 3.1]. We can proceed by induction over n. For $n = N$ the claim trivially holds true, and assume it holds for $0 \le n+1 \le N-1$. Define the events

$$A(n) := \left\{ Z_{t_n} \ge E\left[Z_{\tau_{n+1}} | \mathcal{F}^{\mathbf{X}}_{t_n} \right] \right\} \quad \text{and} \quad A(n,K) := \left\{ Z_{t_n} \ge \psi^{n,K}(\mathbf{X}|_{[0,t_n]}) \right\}.$$

By definition we can write

$$\tau_n^K = t_n 1_{A(n,K)} + \tau_{n+1}^K 1_{A(n,K)^C}, \quad \tau_n = t_n 1_{A(n)} + \tau_{n+1} 1_{A(n)^C}.$$

Using this, it is possible to check that

$$\begin{aligned} E[Z_{\tau_n^K} - Z_{\tau_n} | \mathcal{F}^{\mathbf{X}}_{t_n}] = (Z_{t_n} - &E[Z_{\tau_{n+1}} | \mathcal{F}^{\mathbf{X}}_{t_n}])(1_{A(n,K)} - 1_{A(n)}) \\ &+ E[Z_{\tau_{n+1}^K} - Z_{\tau_{n+1}} | \mathcal{F}^{\mathbf{X}}_{t_n}] 1_{A(n,K)^C}. \end{aligned}$$

The second term converges by induction hypothesis, and we only need to show

$$L_n^K := (Z_{t_n} - E[Z_{\tau_{n+1}} | \mathcal{F}^{\mathbf{X}}_{t_n}])(1_{A(n,K)} - 1_{A(n)}) \xrightarrow{K\to\infty} 0, \text{ in } L^2.$$

Now on $A(n,K) \cap A(n)$ and $A(n,K)^c \cap A(n)^c$ we clearly have $L_n^K = 0$. Moreover

$$1_{A(n,K)^c \cap A(n)} |L_n^K| \le 1_{A(n,K)^c \cap A(n)} |\psi^{n,K}(\mathbf{X}|_{[0,t_n]}) - E[Z_{\tau_{n+1}} | \mathcal{F}^{\mathbf{X}}_{t_n}]|,$$

since $\psi^{n,K}(\mathbf{X}|_{[0,t_n]}) > Z_{t_n} \ge E[Z_{\tau_{n+1}} | \mathcal{F}^{\mathbf{X}}_{t_n}]$ on $A(n,K)^c \cap A(n)$. Similarly, one can show

$$1_{A(n,K) \cap A(n)^c} |L_n^K| \le 1_{A(n,K) \cap A(n)^c} |\psi^{n,K}(\mathbf{X}|_{[0,t_n]}) - E[Z_{\tau_{n+1}} | \mathcal{F}^{\mathbf{X}}_{t_n}]|,$$

and thus

$$|L_n^K| \le |\psi^{n,K}(\mathbf{X}|_{[0,t_n]}) - E[Z_{\tau_{n+1}} | \mathcal{F}^{\mathbf{X}}_{t_n}]|. \tag{34}$$

Notice that $\psi^{n,K}$ is the orthogonal projection of the L^2 random variable $Z_{\tau_{n+1}^K}$ onto the subspace $\{\langle \mathbf{X}^{\le K}_{0,t_n}, l \rangle : l \in \mathcal{W}^{d+1}\}$, and similarly denote by $\hat{\psi}^{n,K}$ the orthogonal projection of $Z_{\tau_{n+1}}$ to the same space. Then we have

$$\begin{aligned} \|L_n^K\|_{L^2} &\le \|\psi^{n,K}(\mathbf{X}|_{[0,t_n]}) - \hat{\psi}^{n,K}(\mathbf{X}|_{[0,t_n]})\|_{L^2} + \|\hat{\psi}^{n,K}(\mathbf{X}|_{[0,t_n]}) - E[Z_{\tau_{n+1}} | \mathcal{F}^{\mathbf{X}}_{t_n}]\|_{L^2} \\ &\le \|E[Z_{\tau_{n+1}^K} | \mathcal{F}^{\mathbf{X}}_{t_n}] - E[Z_{\tau_{n+1}} | \mathcal{F}^{\mathbf{X}}_{t_n}]\|_{L^2} + \|\hat{\psi}^{n,K}(\mathbf{X}|_{[0,t_n]}) - E[Z_{\tau_{n+1}} | \mathcal{F}^{\mathbf{X}}_{t_n}]\|_{L^2}. \end{aligned}$$

Now the first term converges by the induction hypothesis. For the second term, the conditional expectation of the L^2 random variable $Z_{\tau_{n+1}}$ is nothing else than the orthogonal projection onto the space $L^2(\mathcal{F}^{\mathbf{X}}_{t_n})$. But by Theorem 2.6, for any $\epsilon > 0$ we can find $\phi \in L^{\lambda}_{\text{Sig}}$, such that $\|\phi(\mathbf{X}|_{[0,t_n]}) - Z_{\tau_{n+1}}\|_{L^2} \le \epsilon$. For K large enough we have $\phi(\mathbf{X}|_{[0,t_n]}) \in \{\langle \mathbf{X}^{\le K}_{0,t_n}, l \rangle : l \in \mathcal{W}^{d+1}\}$, and thus

$$\|\hat{\psi}^{n,K}(\mathbf{X}|_{[0,t_n]}) - E[Z_{\tau_{n+1}} | \mathcal{F}^{\mathbf{X}}_{t_n}]\|_{L^2} \le \|\hat{\psi}^{n,K}(\mathbf{X}|_{[0,t_n]}) - Z_{\tau_{n+1}}\|_{L^2} \le \epsilon,$$

since $\hat{\psi}^{n,K}$ is such that the distance is minimal. □

Proof of Proposition 3.3 First, we can consider the sequence of stopping times $(\tau_n^{K,J})$ as defined in (21). One can then rewrite exactly the same proof of Proposition 3.2 for $Z_{\tau_n^{K,J}}$ instead, and at the equation (39), we get

$$\begin{aligned} L_n^K &\leq |\psi^{n,K}(\mathbf{X}|_{[0,t_n]}(J)) - E[Z_{\tau_{n+1}}|\mathcal{F}_{t_n}^{\mathbf{X}}]| \\ &\leq |\psi^{n,K}(\mathbf{X}|_{[0,t_n]}) - E[Z_{\tau_{n+1}}|\mathcal{F}_{t_n}^{\mathbf{X}}]| + |\psi^{n,K}(\mathbf{X}|_{[0,t_n]}(J)) - \psi^{n,K}(\mathbf{X}|_{[0,t_n]})|, \end{aligned}$$

where then the first term converges in L^2 due to the same argument as in the proof of Proposition 3.2, and the latter converges by assumption, as $J \to \infty$. It therefore suffices to show that

$$\frac{1}{M}\sum_{i=1}^{M} Z^{(i)}_{\tau_n^{K,J,(i)}} \xrightarrow{M\to\infty} E[Z_{\tau_n^{K,J}}] \text{ a.s.}$$

Now for any $l \in \mathcal{W}^{d+1}$, we can write $l = \lambda_1 w_1 + \cdots + \lambda_D w_D$, where $D = \sum_{k=0}^{K}(d+1)^k$, that is we sum over all possible words of length at most K. One can therefore notice that minimizing $\langle \mathbf{X}^{\leq K}, l\rangle$ over $l \in \mathcal{W}^{d+1}_{\leq K}$, is equivalent to minimizing $\sum_{i=1}^{D}\lambda_i\langle \mathbf{X}^{\leq K}, w_i\rangle$ over all vectors $\lambda \in \mathbb{R}^D$. Defining $e_k(\mathbf{x}) := \langle \mathbf{x}^{\leq K}, w_k\rangle$ for $k = 1,\dots,D$, and setting $X_n := \mathbf{X}|_{[0,t_n]}$, we are exactly in framework of [14, Chapter 3], and the result follows from [14, Theorem 3.2], under the following remark. The authors make the following assumption, denoted by (A2)

$$\sum_j \alpha_j e_j(X_t) = 0 \text{ almost surely implies } \alpha = 0, \forall t \tag{35}$$

for the set of basis-functions, which allows an explicit representation of the coefficient $l^\star$ in (19). Of course, in our framework, such an assumption cannot hold true, as this would correspond to

$$\sum_{l=1}^{D} \alpha_l \langle \mathbf{X}_{0,t}^{<\infty}, w_l\rangle = 0 \text{ a.s.} \implies \alpha_l = 0, \forall l = 1,\dots,D.$$

Since we consider the signature of the time-augmented path (t, X_t), the purely deterministic components of the signature contradict this assumption. However, for a fixed signature level K, we can always discard linear-dependent (in the sense of (40)) components of the signature, that is minimize over the basis-functions

$$\{\tilde{e}_1,\dots,\tilde{e}_{\tilde{D}}\} \subset \{e_1,\dots,e_D\} \text{ s.t. (A2) holds },$$

for the largest possible $\tilde{D} \leq D$. The resulting least-square problem (19) over $\mathbb{R}^{\tilde{D}}$, with respect to $\{\tilde{e}_1,\dots,\tilde{e}_{\tilde{D}}\}$, has an explicit representation of the solution, and since the two sets of basis-functions generate the same subspace of L^2, the explicit solution is also optimal for the original problem. Thus, for a fixed level K, we can proceed with the reduced set of basis-functions, for which the assumption (A2) holds by definition, and we can apply [14, Theorem 3.2]. □

A.2 Proofs in Section 3.3

Proof of Proposition 3.6 The existence of a minimizer is proved in Lemma A.1. We can find the discrete Doob-martingale $M^{\star,N}$ and write

$$y_0^N = E\left[\max_{0\leq n\leq N}\left(Z_{t_n} - M_{t_n}^{\star,N}\right)\right].$$

Define the continuous-time, $(\mathcal{F}_t^{\mathbf{X}})-$martingale $M_t := E[M_T^{\star,N}|\mathcal{F}_t^{\mathbf{X}}]$, and notice that $M_{t_n} = M_{t_n}^{\star,N}$. An application of the martingale approximation in Theorem 3.5 shows that for all $\epsilon > 0$, there exist $l^\epsilon = (l^{i,\epsilon})_{i=1}^m$ in $\left(\mathcal{W}^{d+1}\right)^m$, such that $E\left[\max_{0\le n\le N}\left(M_{t_n}^{\star,N} - M_{t_n}^{l^\epsilon}\right)\right] \le \epsilon$. Thus we have

$$y_0^N = E\left[\max_{0\le n\le N}\left(Z_{t_n} - M_{t_n}^{\star,N}\right)\right] \ge E\left[\max_{0\le n\le N}\left(Z_{t_n} - M_{t_n}^{l^\epsilon}\right)\right] - \epsilon.$$

Now since $y_0^{K,N} \ge y_0^N$, we can find K large enough, such that

$$\begin{aligned} 0 \le y_0^{K,N} - y_0^N &\le \inf_{l\in\left(\mathcal{W}_{\le K}^{d+1}\right)^m} E\left[\max_{0\le n\le N}\left(Z_{t_n} - M_{t_n}^{l}\right)\right] - E\left[\max_{0\le n\le N}\left(Z_{t_n} - M_{t_n}^{l^\epsilon}\right)\right] + \epsilon \\ &\le \epsilon, \end{aligned}$$

where the last inequality follows from that fact that $l^\epsilon \in \left(\mathcal{W}_{\le K}^{d+1}\right)^m$ for K large enough. □

Lemma A.1 *The minimization problem*

$$y_0^{K,N} = \inf_{l\in\left(\mathcal{W}_{\le K}^{d+1}\right)^m} E\left[\max_{0\le n\le N}\left(Z_{t_n} - M_{t_n}^{l}\right)\right]$$

has a solution.

Proof First notice that $l \mapsto E\left[\max_{0\le n\le N}\left(Z_{t_n} - M_{t_n}^{l}\right)\right]$ is convex. Then we have

$$\begin{aligned} E\left[\max_{0\le n\le N}\left(Z_{t_n} - M_{t_n}^{l}\right)\right] &\ge E\left[\max\left(Z_T - M_T^l, 0\right)\right] \\ &= \frac{1}{2}E\left[Z_T - M_T^l + \left|M_T^l - Z_T\right|\right] \\ &\ge \frac{1}{2}E\left[\left|M_T^l\right|\right] + E[\max(-Z_T, 0)], \end{aligned}$$

where the equality in the middle uses $\max(A-B,0) = \frac{1}{2}\left(A - B + |B-A|\right)$. Now for any word $l = \lambda_1 w_1 + \cdots + \lambda_n w_n$, we set $|l| = \sum_{i=1}^n |\lambda_i|$, and notice that

$$\frac{1}{2}E\left[|M_T|\right] = \frac{1}{2}|l|E[|M_T^{l/|l|}|] \ge \frac{|l|}{2} \inf_{\widehat{l}\in\left(\mathcal{W}_{\le K}^{d+1}\right)^m, |\widehat{l}|=1} E[|M_T^{\widehat{l}}|]. \tag{36}$$

Since $\widehat{l} \mapsto E[|M_T^{\widehat{l}}|]$ is continuous and the set $\{\widehat{l} \in \left(\mathcal{W}_{\le K}^{d+1}\right)^m : |\widehat{l}| = 1\}$ is compact, the minimum on the right hand-side of (41) is attained. Assume now that $\inf_{\widehat{l}\in\left(\mathcal{W}_{\le K}^{d+1}\right)^m, |\widehat{l}|=1} E[|M_T^{\widehat{l}}|] = 0$. Then there exists an $\widehat{l}^\star$ with $|\widehat{l}^\star| = 1$ and $|M_T^{\widehat{l}^\star}| = 0$ almost surely. But this in particular implies that

$$\langle \mathbf{X}_{0,s}^{<\infty}, \widehat{l}^\star\rangle = 0, \quad \text{for almost every } s \in [0,T] \text{ almost surely.}$$

But this is only possible if $\widehat{l}^\star = 0$, contradicting the fact that $|\widehat{l}^\star| = 1$. Hence the infimum (41) is positive and we can conclude that the function

$$l \mapsto E\left[\max_{0\le n\le N}\left(Z_{t_n} - M_{t_n}^{l}\right)\right] \xrightarrow{|l|\to\infty} \infty,$$

which implies the existence of the minimizer. □

Finally, in order to prove Proposition 3.8, we quickly introduce the general idea of sample average approximation (SAA), for which we refer to [34, Chapter 6] for details. Assume $\mathcal{X}$ is a closed and convex subset of $\mathbb{R}^N$ and ξ is a random vector in $\mathbb{R}^d$ for some $d, N \in \mathbb{N}$, and F is some function $F : \mathbb{R}^N \times \mathbb{R}^d \to \mathbb{R}$. We are interested in approximating the stochastic programming problem

$$y_0 = \min_{x \in \mathcal{X}} E[F(x, \xi)]. \tag{37}$$

Define the sample average function $F^M(x) = \frac{1}{M}\sum_{j=1}^M F(x, \xi^j)$, where $\xi^j, j = 1, \ldots, M$ are i.i.d samples of the random vector ξ. The sample average approximation of y_0 is then given by

$$y_0^M = \min_{x \in \mathcal{X}} F^M(x). \tag{38}$$

The following result provides sufficient conditions for the convergence $y_0^M \xrightarrow{M\to\infty} y_0$, and a more general version can be found in [34, Chapter 6 Theorem 4].

Theorem A.2 *Suppose that*

(1) F is measurable and $x \mapsto F(x, \xi)$ is lower semicontinuous for all $\xi \in \mathbb{R}^d$,

(2) $x \mapsto F(x, \xi)$ is convex for almost every ξ,

(3) $\mathcal{X}$ is closed and convex,

(4) $f(x) := E[F(x, \xi)]$ is lower semicontinuous and $f(x) < \infty$ for all $x \in \mathcal{X}$,

(5) the set S of optimal solutions to (42) *is non-empty and bounded.*

Then $y_0^M \xrightarrow{M\to\infty} y_0$.

Using the notation of Section 3.3, for some fixed dimension m of the Brownian motion W, number of exercise dates N, number discretization points J and signature truncation level K, we can define the random vector

$$\xi := (Z_{t_0}, \ldots, Z_{t_N}, \tilde{M}^1_{t_0}, \ldots, \tilde{M}^1_{t_N}, \ldots, \tilde{M}^D_{t_0}, \ldots, \tilde{M}^D_{t_N}) \in \mathbb{R}^{(D+1)\cdot(N+1)\cdot m},$$

where D denotes the number of entries of the signature up to level K. Using the notation $M_t^{\beta,J} = \sum_{j=1}^D \beta_j M_t^{w_j,J}$ for any $\beta \in \mathbb{R}^D$, we can consider the minimization problem

$$y_0^{K,J,N} = \inf_{\beta \in \mathbb{R}^{D\cdot m}} E\left[\max_{0\le n\le N}\left(Z_{t_n} - \tilde{M}^{\beta,J}_{t_n}\right)\right] = \min_{x\in\mathcal{X}} E[F(x,\xi)],$$

where $\mathcal{X} = \mathbb{R}^{D\cdot m}$ and $d = (D+1)(N+1)m$.

Proof of Proposition 3.8 First, it is possible to rewrite the proof of Lemma A.1 for F^M instead of the expectation when M is large enough, to show that there exists a minimizer $l^\star$ to (27). Moreover, the minimization problem $y_0^{K,N,J}$ is nothing else than (24), where the continuous martingales M^β are replaced by the Euler approximations $M^{l,J}$ on the grid $(s_j)_{j=0}^J$. By the same reasoning as in

Proposition 3.6 we can see that $0 \leq y_0^{K,N,J} - y_0^N \longrightarrow 0$ as $K, J \to \infty$. Applying the triangle inequality yields

$$|y_0^{K,N,J,M} - y_0^N| \leq |y_0^{K,N,J} - y_0^N| + |y_0^{K,N,J,M} - y_0^{K,N,J}|,$$

and thus we are left with the convergence of the second term. But this can be deduced from the last Theorem, if we can show that (1)-(5) hold true for our F. Clearly F is measurable and it is easy to see that $x \to F(x,\xi)$ is continuous and convex, thus (1) and (2) readily follow. Moreover (3) holds true, and for $x_1, x_2 \in \mathcal{X}$ we have

$$|f(x_1) - f(x_2)| \lesssim |x_1 - x_2| E[|\tilde{M}_T^{(1,\dots,1)}|^2]$$

using Doobs-inequality. Since $M^l \in L^2$ for all l and $\|Z\|_\infty \in L^1$, (4) follows. Finally, non-emptyness of S follows from Lemma A.1, and the proof of the latter reveals that $E[F(x,\xi)] \xrightarrow{|x|\to\infty} \infty$, and thus S must be bounded, which finishes the proof. □

A Technical details Section 3

A.1 Proofs in Section 3.2

Proof of Proposition 3.2 The proof is based on the same ideas as the proof in [14, Theorem 3.1]. We can proceed by induction over n. For $n = N$ the claim trivially holds true, and assume it holds for $0 \leq n+1 \leq N-1$. Define the events

$$A(n) := \left\{ Z_{t_n} \geq E\left[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}^{\mathbf{X}} \right] \right\} \quad \text{and} \quad A(n,K) := \left\{ Z_{t_n} \geq \psi^{n,K}(\mathbf{X}|_{[0,t_n]}) \right\}.$$

By definition we can write

$$\tau_n^K = t_n 1_{A(n,K)} + \tau_{n+1}^K 1_{A(n,K)^C}, \quad \tau_n = t_n 1_{A(n)} + \tau_{n+1} 1_{A(n)^C}.$$

Using this, it is possible to check that

$$\begin{aligned} E[Z_{\tau_n^K} - Z_{\tau_n} | \mathcal{F}_{t_n}^{\mathbf{X}}] = (Z_{t_n} - &E[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}^{\mathbf{X}}])(1_{A(n,K)} - 1_{A(n)}) \\ &+ E[Z_{\tau_{n+1}^K} - Z_{\tau_{n+1}} | \mathcal{F}_{t_n}^{\mathbf{X}}] 1_{A(n,K)^C}. \end{aligned}$$

The second term converges by induction hypothesis, and we only need to show

$$L_n^K := (Z_{t_n} - E[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}^{\mathbf{X}}])(1_{A(n,K)} - 1_{A(n)}) \xrightarrow{K\to\infty} 0, \text{ in } L^2.$$

Now on $A(n,K) \cap A(n)$ and $A(n,K)^c \cap A(n)^c$ we clearly have $L_n^K = 0$. Moreover

$$1_{A(n,K)^c \cap A(n)} |L_n^K| \leq 1_{A(n,K)^c \cap A(n)} |\psi^{n,K}(\mathbf{X}|_{[0,t_n]}) - E[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}^{\mathbf{X}}]|,$$

since $\psi^{n,K}(\mathbf{X}|_{[0,t_n]}) > Z_{t_n} \geq E[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}^{\mathbf{X}}]$ on $A(n,K)^c \cap A(n)$. Similarly, one can show

$$1_{A(n,K) \cap A(n)^c} |L_n^K| \leq 1_{A(n,K) \cap A(n)^c} |\psi^{n,K}(\mathbf{X}|_{[0,t_n]}) - E[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}^{\mathbf{X}}]|,$$

and thus

$$|L_n^K| \leq |\psi^{n,K}(\mathbf{X}|_{[0,t_n]}) - E[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}^{\mathbf{X}}]|. \tag{39}$$

Notice that $\psi^{n,K}$ is the orthogonal projection of the L^2 random variable $Z_{\tau_{n+1}^K}$ onto the subspace $\{\langle \mathbf{X}_{0,t_n}^{\le K}, l\rangle : l \in \mathcal{W}^{d+1}\}$, and similarly denote by $\hat{\psi}^{n,K}$ the orthogonal projection of $Z_{\tau_{n+1}}$ to the same space. Then we have

$$\begin{aligned}\|L_n^K\|_{L^2} &\le \|\psi^{n,K}(\mathbf{X}|_{[0,t_n]}) - \hat{\psi}^{n,K}(\mathbf{X}|_{[0,t_n]})\|_{L^2} + \|\hat{\psi}^{n,K}(\mathbf{X}|_{[0,t_n]}) - E[Z_{\tau_{n+1}}|\mathcal{F}_{t_n}^{\mathbf{X}}]\|_{L^2} \\ &\le \|E[Z_{\tau_{n+1}^K}|\mathcal{F}_{t_n}^{\mathbf{X}}] - E[Z_{\tau_{n+1}}|\mathcal{F}_{t_n}^{\mathbf{X}}]\|_{L^2} + \|\hat{\psi}^{n,K}(\mathbf{X}|_{[0,t_n]}) - E[Z_{\tau_{n+1}}|\mathcal{F}_{t_n}^{\mathbf{X}}]\|_{L^2}.\end{aligned}$$

Now the first term converges by the induction hypothesis. For the second term, the conditional expectation of the L^2 random variable $Z_{\tau_{n+1}}$ is nothing else than the orthogonal projection onto the space $L^2(\mathcal{F}_{t_n}^{\mathbf{X}})$. But by Theorem 2.6, for any $\epsilon > 0$ we can find $\phi \in L_{\text{Sig}}^{\lambda}$, such that $\|\phi(\mathbf{X}|_{[0,t_n]}) - Z_{\tau_{n+1}}\|_{L^2} \le \epsilon$. For K large enough we have $\phi(\mathbf{X}|_{[0,t_n]}) \in \{\langle \mathbf{X}_{0,t_n}^{\le K}, l\rangle : l \in \mathcal{W}^{d+1}\}$, and thus

$$\|\hat{\psi}^{n,K}(\mathbf{X}|_{[0,t_n]}) - E[Z_{\tau_{n+1}}|\mathcal{F}_{t_n}^{\mathbf{X}}]\|_{L^2} \le \|\hat{\psi}^{n,K}(\mathbf{X}|_{[0,t_n]}) - Z_{\tau_{n+1}}\|_{L^2} \le \epsilon,$$

since $\hat{\psi}^{n,K}$ is such that the distance is minimal. □

Proof of Proposition 3.3 First, we can consider the sequence of stopping times $(\tau_n^{K,J})$ as defined in (21). One can then rewrite exactly the same proof of Proposition 3.2 for $Z_{\tau_n^{K,J}}$ instead, and at the equation (39), we get

$$\begin{aligned}L_n^K &\le |\psi^{n,K}(\mathbf{X}|_{[0,t_n]}(J)) - E[Z_{\tau_{n+1}}|\mathcal{F}_{t_n}^{\mathbf{X}}]| \\ &\le |\psi^{n,K}(\mathbf{X}|_{[0,t_n]}) - E[Z_{\tau_{n+1}}|\mathcal{F}_{t_n}^{\mathbf{X}}]| + |\psi^{n,K}(\mathbf{X}|_{[0,t_n]}(J)) - \psi^{n,K}(\mathbf{X}|_{[0,t_n]})|,\end{aligned}$$

where then the first term converges in L^2 due to the same argument as in the proof of Proposition 3.2, and the latter converges by assumption, as $J \to \infty$. It therefore suffices to show that

$$\frac{1}{M}\sum_{i=1}^{M} Z_{\tau_n^{K,J,(i)}}^{(i)} \xrightarrow{M\to\infty} E[Z_{\tau_n^{K,J}}] \text{ a.s.}$$

Now for any $l \in \mathcal{W}^{d+1}$, we can write $l = \lambda_1 w_1 + \cdots + \lambda_D w_D$, where $D = \sum_{k=0}^{K}(d+1)^k$, that is we sum over all possible words of length at most K. One can therefore notice that minimizing $\langle \mathbf{X}^{\le K}, l\rangle$ over $l \in \mathcal{W}_{\le K}^{d+1}$, is equivalent to minimizing $\sum_{i=1}^{D}\lambda_i\langle \mathbf{X}^{\le K}, w_i\rangle$ over all vectors $\lambda \in \mathbb{R}^D$. Defining $e_k(\mathbf{x}) := \langle \mathbf{x}^{\le K}, w_k\rangle$ for $k = 1, \ldots, D$, and setting $X_n := \mathbf{X}|_{[0,t_n]}$, we are exactly in framework of [14, Chapter 3], and the result follows from [14, Theorem 3.2], under the following remark. The authors make the following assumption, denoted by (A2)

$$\sum_j \alpha_j e_j(X_t) = 0 \text{ almost surely implies } \alpha = 0, \forall t \tag{40}$$

for the set of basis-functions, which allows an explicit representation of the coefficient $l^\star$ in (19). Of course, in our framework, such an assumption cannot hold true, as this would correspond to

$$\sum_{l=1}^{D}\alpha_l\langle \mathbf{X}_{0,t}^{<\infty}, w_l\rangle = 0 \text{ a.s.} \implies \alpha_l = 0, \forall l = 1, \ldots, D.$$

Since we consider the signature of the time-augmented path (t, X_t), the purely deterministic components of the signature contradict this assumption. However, for a fixed signature level K, we can always discard linear-dependent (in the sense of (40)) components of the signature, that is minimize over the basis-functions

$$\{\tilde{e}_1, \ldots, \tilde{e}_{\tilde{D}}\} \subset \{e_1, \ldots, e_D\} \text{ s.t. (A2) holds },$$

for the largest possible $\tilde{D} \leq D$. The resulting least-square problem (19) over $\mathbb{R}^{\tilde{D}}$, with respect to $\{\tilde{e}_1, \ldots, \tilde{e}_{\tilde{D}}\}$, has an explicit representation of the solution, and since the two sets of basis-functions generate the same subspace of L^2, the explicit solution is also optimal for the original problem. Thus, for a fixed level K, we can proceed with the reduced set of basis-functions, for which the assumption (A2) holds by definition, and we can apply [14, Theorem 3.2]. □

A.2 Proofs in Section 3.3

Proof of Proposition 3.6 The existence of a minimizer is proved in Lemma A.1. We can find the discrete Doob-martingale $M^{\star,N}$ and write

$$y_0^N = E\left[\max_{0\leq n\leq N}\left(Z_{t_n} - M_{t_n}^{\star,N}\right)\right].$$

Define the continuous-time, $(\mathcal{F}_t^{\mathbf{X}})-$martingale $M_t := E[M_T^{\star,N}|\mathcal{F}_t^{\mathbf{X}}]$, and notice that $M_{t_n} = M_{t_n}^{\star,N}$. An application of the martingale approximation in Theorem 3.5 shows that for all $\epsilon > 0$, there exist $l^\epsilon = (l^{i,\epsilon})_{i=1}^m$ in $\left(\mathcal{W}^{d+1}\right)^m$, such that $E\left[\max_{0\leq n\leq N}\left(M_{t_n}^{\star,N} - M_{t_n}^{l^\epsilon}\right)\right] \leq \epsilon$. Thus we have

$$y_0^N = E\left[\max_{0\leq n\leq N}\left(Z_{t_n} - M_{t_n}^{\star,N}\right)\right] \geq E\left[\max_{0\leq n\leq N}\left(Z_{t_n} - M_{t_n}^{l^\epsilon}\right)\right] - \epsilon.$$

Now since $y_0^{K,N} \geq y_0^N$, we can find K large enough, such that

$$\begin{aligned} 0 \leq y_0^{K,N} - y_0^N &\leq \inf_{l\in\left(\mathcal{W}_{\leq K}^{d+1}\right)^m} E\left[\max_{0\leq n\leq N}\left(Z_{t_n} - M_{t_n}^{l}\right)\right] - E\left[\max_{0\leq n\leq N}\left(Z_{t_n} - M_{t_n}^{l^\epsilon}\right)\right] + \epsilon \\ &\leq \epsilon, \end{aligned}$$

where the last inequality follows from that fact that $l^\epsilon \in \left(\mathcal{W}_{\leq K}^{d+1}\right)^m$ for K large enough. □

Lemma A.1 *The minimization problem*

$$y_0^{K,N} = \inf_{l\in\left(\mathcal{W}_{\leq K}^{d+1}\right)^m} E\left[\max_{0\leq n\leq N}\left(Z_{t_n} - M_{t_n}^{l}\right)\right]$$

has a solution.

Proof First notice that $l \mapsto E\left[\max_{0\leq n\leq N}\left(Z_{t_n} - M_{t_n}^{l}\right)\right]$ is convex. Then we have

$$\begin{aligned} E\left[\max_{0\leq n\leq N}\left(Z_{t_n} - M_{t_n}^{l}\right)\right] &\geq E\left[\max\left(Z_T - M_T^l, 0\right)\right] \\ &= \frac{1}{2}E\left[Z_T - M_T^l + \left|M_T^l - Z_T\right|\right] \\ &\geq \frac{1}{2}E\left[\left|M_T^l\right|\right] + E[\max(-Z_T, 0)], \end{aligned}$$

where the equality in the middle uses $\max(A-B,0) = \frac{1}{2}(A - B + |B - A|)$. Now for any word $l = \lambda_1 w_1 + \cdots + \lambda_n w_n$, we set $|l| = \sum_{i=1}^n |\lambda_i|$, and notice that

$$\frac{1}{2}E\left[|M_T|\right] = \frac{1}{2}|l|E[|M_T^{l/|l|}|] \geq \frac{|l|}{2}\inf_{\hat{l}\in\left(\mathcal{W}_{\leq K}^{d+1}\right)^m, |\hat{l}|=1} E[|M_T^{\hat{l}}|]. \tag{41}$$

Since $\widehat{l} \mapsto E[|M_T^{\widehat{l}}|]$ is continuous and the set $\{\widehat{l} \in \left(\mathcal{W}_{\leq K}^{d+1}\right)^m : |\widehat{l}| = 1\}$ is compact, the minimum on the right hand-side of (41) is attained. Assume now that $\inf_{\widehat{l} \in \left(\mathcal{W}_{\leq K}^{d+1}\right)^m, |\widehat{l}|=1} E[|M_T^{\widehat{l}}|] = 0$. Then there exists an $\widehat{l^\star}$ with $|\widehat{l^\star}| = 1$ and $|M_T^{\widehat{l^\star}}| = 0$ almost surely. But this in particular implies that

$$\langle \mathbf{X}_{0,s}^{<\infty}, \widehat{l^\star} \rangle = 0, \quad \text{for almost every } s \in [0,T] \text{ almost surely.}$$

But this is only possible if $\widehat{l^\star} = 0$, contradicting the fact that $|\widehat{l^\star}| = 1$. Hence the infimum (41) is positive and we can conclude that the function

$$l \mapsto E\left[\max_{0 \leq n \leq N} \left(Z_{t_n} - M_{t_n}^l\right)\right] \xrightarrow{|l| \to \infty} \infty,$$

which implies the existence of the minimizer. □

Finally, in order to prove Proposition 3.8, we quickly introduce the general idea of sample average approximation (SAA), for which we refer to [34, Chapter 6] for details. Assume $\mathcal{X}$ is a closed and convex subset of $\mathbb{R}^N$ and ξ is a random vector in $\mathbb{R}^d$ for some $d, N \in \mathbb{N}$, and F is some function $F : \mathbb{R}^N \times \mathbb{R}^d \to \mathbb{R}$. We are interested in approximating the stochastic programming problem

$$y_0 = \min_{x \in \mathcal{X}} E[F(x, \xi)]. \tag{42}$$

Define the sample average function $F^M(x) = \frac{1}{M} \sum_{j=1}^M F(x, \xi^j)$, where $\xi^j, j = 1, \ldots, M$ are i.i.d samples of the random vector ξ. The sample average approximation of y_0 is then given by

$$y_0^M = \min_{x \in \mathcal{X}} F^M(x). \tag{43}$$

The following result provides sufficient conditions for the convergence $y_0^M \xrightarrow{M \to \infty} y_0$, and a more general version can be found in [34, Chapter 6 Theorem 4].

Theorem A.2 *Suppose that*

(1) F is measurable and $x \mapsto F(x, \xi)$ is lower semicontinuous for all $\xi \in \mathbb{R}^d$,

(2) $x \mapsto F(x, \xi)$ is convex for almost every ξ,

(3) $\mathcal{X}$ is closed and convex,

(4) $f(x) := E[F(x, \xi)]$ is lower semicontinuous and $f(x) < \infty$ for all $x \in \mathcal{X}$,

(5) the set S of optimal solutions to (42) *is non-empty and bounded.*

Then $y_0^M \xrightarrow{M \to \infty} y_0$.

Using the notation of Section 3.3, for some fixed dimension m of the Brownian motion W, number of exercise dates N, number discretization points J and signature truncation level K, we can define the random vector

$$\xi := (Z_{t_0}, \ldots, Z_{t_N}, \tilde{M}_{t_0}^1, \ldots, \tilde{M}_{t_N}^1, \ldots, \tilde{M}_{t_0}^D, \ldots, \tilde{M}_{t_N}^D) \in \mathbb{R}^{(D+1) \cdot (N+1) \cdot m},$$

where D denotes the number of entries of the signature up to level K. Using the notation $M_t^{\beta,J} = \sum_{j=1}^D \beta_j M_t^{w_j,J}$ for any $\beta \in \mathbb{R}^D$, we can consider the minimization problem

$$y_0^{K,J,N} = \inf_{\beta \in \mathbb{R}^{D\cdot m}} E\left[\max_{0\le n\le N}\left(Z_{t_n} - \tilde{M}_{t_n}^{\beta,J}\right)\right] = \min_{x\in\mathcal{X}} E[F(x,\xi)],$$

where $\mathcal{X} = \mathbb{R}^{D\cdot m}$ and $d = (D+1)(N+1)m$.

Proof of Proposition 3.8 First, it is possible to rewrite the proof of Lemma A.1 for F^M instead of the expectation when M is large enough, to show that there exists a minimizer $l^\star$ to (27). Moreover, the minimization problem $y_0^{K,N,J}$ is nothing else than (24), where the continuous martingales M^β are replaced by the Euler approximations $M^{l,J}$ on the grid $(s_j)_{j=0}^J$. By the same reasoning as in Proposition 3.6 we can see that $0 \le y_0^{K,N,J} - y_0^N \longrightarrow 0$ as $K, J \to \infty$. Applying the triangle inequality yields

$$|y_0^{K,N,J,M} - y_0^N| \le |y_0^{K,N,J} - y_0^N| + |y_0^{K,N,J,M} - y_0^{K,N,J}|,$$

and thus we are left with the convergence of the second term. But this can be deduced from the last Theorem, if we can show that (1)-(5) hold true for our F. Clearly F is measurable and it is easy to see that $x \to F(x,\xi)$ is continuous and convex, thus (1) and (2) readily follow. Moreover (3) holds true, and for $x_1, x_2 \in \mathcal{X}$ we have

$$|f(x_1) - f(x_2)| \lesssim |x_1 - x_2| E[|\tilde{M}_T^{(1,\dots,1)}|^2]$$

using Doobs-inequality. Since $M^l \in L^2$ for all l and $\|Z\|_\infty \in L^1$, (4) follows. Finally, non-emptyness of S follows from Lemma A.1, and the proof of the latter reveals that $E[F(x,\xi)] \xrightarrow{|x|\to\infty} \infty$, and thus S must be bounded, which finishs the proof. □

References

[1] Leif Andersen and Mark Broadie. Primal-dual simulation algorithm for pricing multidimensional American options. *Management Science*, 50(9):1222–1234, 2004.

[2] Christian Bayer and Simon Breneis. Efficient option pricing in the rough Heston model using weak simulation schemes. *arXiv preprint arXiv:2310.04146*, 2023.

[3] Christian Bayer, Peter Friz, and Jim Gatheral. Pricing under rough volatility. *Quantitative Finance*, 16(6):887–904, 2016.

[4] Christian Bayer, Paul P Hager, Sebastian Riedel, and John Schoenmakers. Optimal stopping with signatures. *The Annals of Applied Probability*, 33(1):238–273, 2023.

[5] Christian Bayer, Raúl Tempone, and Sören Wolfers. Pricing American options by exercise rate optimization. *Quantitative Finance*, 20(11):1749–1760, 2020.

[6] Sebastian Becker, Patrick Cheridito, and Arnulf Jentzen. Deep optimal stopping. *The Journal of Machine Learning Research*, 20(1):2712–2736, 2019.

[7] Denis Belomestny, Christian Bender, and John Schoenmakers. Solving optimal stopping problems via randomization and empirical dual optimization. *Mathematics of Operations Research*, 48(3):1454–1480, 2023.

[8] Denis Belomestny and John Schoenmakers. From optimal martingales to randomized dual optimal stopping. *Quantitative Finance*, 23(7-8):1099–1113, 2023.

[9] Horatio Boedihardjo, Xi Geng, Terry Lyons, and Danyu Yang. The signature of a rough path: uniqueness. *Advances in Mathematics*, 293:720–737, 2016.

[10] Vladimir Igorevich Bogachev and Maria Aparecida Soares Ruas. *Measure theory*, volume 1. Springer, 2007.

[11] Jean-Philippe Bouchaud, Julius Bonart, Jonathan Donier, and Martin Gould. *Trades, quotes and prices: financial markets under the microscope*. Cambridge University Press, 2018.

[12] Philippe Carmona and Laure Coutin. Fractional Brownian motion and the Markov property. *Electronic Communications in Probability [electronic only]*, 3:95–107, 1998.

[13] Ilya Chevyrev and Harald Oberhauser. Signature moments to characterize laws of stochastic processes. *The Journal of Machine Learning Research*, 23(1):7928–7969, 2022.

[14] Emmanuelle Clément, Damien Lamberton, and Philip Protter. An analysis of a least squares regression method for American option pricing. *Finance and stochastics*, 6(4):449, 2002.

[15] Rama Cont and David-Antoine Fournié. Functional Itô calculus and stochastic integral representation of martingales. *The Annals of Applied Probability*, 41(1):109–133, 2013.

[16] Christa Cuchiero, Philipp Schmocker, and Josef Teichmann. Global universal approximation of functional input maps on weighted spaces. *arXiv preprint arXiv:2306.03303*, 2023.

[17] Christa Cuchiero and Josef Teichmann. Generalized Feller processes and Markovian lifts of stochastic Volterra processes: the affine case. *Journal of evolution equations*, 20(4):1301–1348, 2020.

[18] Vijay V Desai, Vivek F Farias, and Ciamac C Moallemi. Pathwise optimization for optimal stopping problems. *Management Science*, 58(12):2292–2308, 2012.

[19] Bruno Dupire. Functional Itô calculus. *Quantitative Finance*, 19(5):721–729, 2019.

[20] Peter K Friz, Paul P Hager, and Nikolas Tapia. Unified signature cumulants and generalized Magnus expansions. In *Forum of Mathematics, Sigma*, volume 10. Cambridge University Press, 2022.

[21] Peter K Friz and Martin Hairer. *A course on rough paths*. Springer, 2020.

[22] Peter K Friz and Nicolas B Victoir. *Multidimensional stochastic processes as rough paths: theory and applications*, volume 120. Cambridge University Press, 2010.

[23] Jim Gatheral, Thibault Jaisson, and Mathieu Rosenbaum. Volatility is rough. *Quantitative Finance*, 18(6):933–949, 2018.

[24] Robin Giles. A generalization of the strict topology. *Transactions of the American Mathematical Society*, 161:467–474, 1971.

[25] Ludovic Goudenege, Andrea Molent, and Antonino Zanette. Machine learning for pricing American options in high-dimensional Markovian and non-Markovian models. *Quantitative Finance*, 20(4):573–591, 2020.

[26] Julien Guyon and Jordan Lekeufack. Volatility is (mostly) path-dependent. *Quantitative Finance*, 23(9):1221–1258, 2023.

[27] Jasdeep Kalsi, Terry Lyons, and Imanol Perez Arribas. Optimal execution with rough path signatures. *SIAM Journal on Financial Mathematics*, 11(2):470–493, 2020.

[28] Ioannis Karatzas and Steven Shreve. *Brownian motion and stochastic calculus*, volume 113. Springer Science & Business Media, 1991.

[29] Francis A Longstaff and Eduardo S Schwartz. Valuing American options by simulation: a simple least-squares approach. *The review of financial studies*, 14(1):113–147, 2001.

[30] Terry J Lyons. Differential equations driven by rough signals. *Revista Matemática Iberoamericana*, 14(2):215–310, 1998.

[31] Yosihiko Ogata. Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical association*, 83(401):9–27, 1988.

[32] Goran Peskir and Albert Shiryaev. *Optimal stopping and free-boundary problems*. Springer, 2006.

[33] Leonard CG Rogers. Monte carlo valuation of American options. *Mathematical Finance*, 12(3):271–286, 2002.

[34] Alexander Shapiro. Monte carlo sampling methods. *Handbooks in operations research and management science*, 10:353–425, 2003.