

## **Pricing American options under rough volatility using deep-signatures and signature-kernels**

Christian Bayer<sup>1</sup>, Luca Pelizzari<sup>1,2</sup>, Jia-Jie Zhu<sup>1</sup>

submitted: January 31, 2025

<sup>1</sup> Weierstrass Institute  
Mohrenstr. 39  
10117 Berlin  
Germany  
E-Mail: christian.bayer@wias-berlin.de  
luca.pelizzari@wias-berlin.de  
jia-jie.zhu@wias-berlin.de

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

No. 3172  
Berlin 2025



---

2020 *Mathematics Subject Classification.* 60G40,60L10,91G20, 91G60.

*Key words and phrases.* Signature, optimal stopping, rough volatility, deep learning, kernel learning.

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). CB and PL acknowledge support from DFG CRC/TRR 388 “Rough Analysis, Stochastic Dynamics and Related Fields”, Project B03.

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](mailto:preprint@wias-berlin.de)  
World Wide Web: <http://www.wias-berlin.de/>

# Pricing American options under rough volatility using deep-signatures and signature-kernels

Christian Bayer, Luca Pelizzari, Jia-Jie Zhu

## Abstract

We extend the signature-based primal and dual solutions to the optimal stopping problem recently introduced in [Bayer et al.: Primal and dual optimal stopping with signatures, to appear in Finance & Stochastics 2025], by integrating deep-signature and signature-kernel learning methodologies. These approaches are designed for non-Markovian frameworks, in particular enabling the pricing of American options under rough volatility. We demonstrate and compare the performance within the popular rough Heston and rough Bergomi models.

## 1 Introduction

Starting with the seminal paper [19], (rough) path signatures [11] have been increasingly recognized as a powerful tool for numerical approximations to solutions of stochastic optimal control problems when the underlying system does not have the Markov property. While the workhorse theoretical methods for stochastic optimal control – dynamic programming / HJB equations on the one hand, Pontryagin maximum principle / BSDEs on the other hand – can, in principle, be formulated and analysed even when the underlying system is not Markovian, this comes at the expense of tractability.

For example, in the Markovian case, we can assume (under mild conditions) that optimal controls are of feedback form, i.e., can be expressed as functions  $\alpha^*(t, X_t)$  of the underlying state variable  $X_t$ . If we do not make the Markovian assumption, we only know that optimal controls  $\alpha_t^*$  are, say, progressively measurable w.r.t. the governing filtration, i.e.,  $\alpha_t^* = \alpha^*(t, X|_{[0,t]})$ , provided that said filtration is generated by a process  $X$ . Similar remarks can be made for the value function as well as the solutions to the corresponding BSDE systems.

From a numerical point of view, this observation increases the computational challenge considerably for solving non-Markovian stochastic optimal control problems. In essence, an approximation problem for a finite-dimensional function, say,  $\alpha^* : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^m$ , is replaced with one for a function on pathspace, say  $\alpha^* : [0, T] \times C([0, T]; \mathbb{R}^d) \rightarrow \mathbb{R}^m$  (with additional measurability constraints).

Kalsi, Lyons and Perez Arribas suggested a general framework for solving stochastic optimal control problems in a model-free way (in particular, without assuming a Markovian problem) for the example of an optimal execution problem, see [19]. In a nutshell, the approach consists of approximating the strategy as well as the value function as linear functionals of the signature of the underlying process. The *linearization* – based on the signature's *universality* – allows them to rephrase the optimal execution problem as a *deterministic* optimization problem in terms of the *expected signature*.

Full linearization as in [19] is only feasible if the unknown function (value function, control, ...) is smooth enough, so as not to require a very deep degree of signature approximation – comparable to approximations with polynomials of high degree in finite dimensional spaces. Otherwise, several other strategies have been considered in the literature:

- Signatures can be interlaced with non-linear transformations of the path, e.g., deep neural networks, to increase the signature's “expressivity”, see, for instance, [35].
- Signatures – or, alternatively, log-signatures to reduce the dimension – can be used as features for other, non-linear approximation methods, e.g., deep neural networks. In the context of optimal stopping, this was first suggested in [5].

In addition, *signature kernels*, see [12], provide a classical kernel learning approach to regression problems in general, and can also be used for stochastic optimal control, in particular. [32] observed that the (non-truncated) signature kernel satisfies a Goursat PDE, and, hence, can be computed numerically, without requiring the (otherwise crucial) truncations of the signature.

*Rough volatility models*, see [4], are a popular class of stochastic volatility models in finance, i.e., with a stock price process following a dynamics  $dS_t = \sqrt{v_t} S_t dZ_t$ , where the stochastic variance process  $v$  is “rough”, e.g., an exponential of fractional Brownian motion with Hurst index  $0 < H < 1/2$ . Crucially, such models  $(S, v)$  do not have the Markov property, so signature methods are ideally suited for computing approximate solutions to the optimal stopping problem in  $S$  – in financial terms: to compute the prices of *American or Bermudan* options.

This paper builds on [6], where adaptations of the classical Longstaff–Schwartz and dual algorithms for Bermudan option pricing based on linear functionals of the signature were introduced and analysed. In this paper, we further extend these algorithms by introducing versions based on non-linear functions of the signature as well as signature-kernel versions. We then apply them to the Bermudan options pricing problem for popular rough volatility models (specifically, the rough Bergomi model [3] and the rough Heston model [15]), and numerically analyse their performance for these models under realistic model parameters.

The goal of this paper is to compute American (more precisely, Bermudan) option prices in the aforementioned rough volatility models, with special emphasis on numerical performance as well as comparisons between the methods outlined before. Specifically, we solve primal and dual formulations of the optimal stopping problem – giving us lower and upper bounds of the option price, respectively – using three different signature methods each:

- 1 linear functionals of the truncated signature;
- 2 deep neural networks applied to truncated signatures;
- 3 linear combinations of the signature kernel.

Compared to the literature (see, for instance, [6] based on linear functionals of the signature alone), we provide considerably sharper bounds for the option price, in the sense of the length of the interval between lower and upper bounds. More specifically, we find that for a realistic rough volatility model (the rough Bergomi model with Hurst index  $H = 0.07$ , as suggested in [3]), we obtain a tight gap of about 1%.

Regarding the comparison between the different methods, no method seems to consistently outperform the others. For the primal problem, all three methods tend to perform very well, with mild advantages for the linear and the signature kernel methods. The dual formulation, however, seems to lead to more difficult approximation and optimization problems, and the neural-network-based method has some advantages. Overall, we see the biggest improvements compared to [6] for the dual formulation.

This conclusion is also supported by various other numerical studies performed, including a study of the dependence of the training error on the number of training samples. We also study the relative

importance of the different components in the signature, specifically in the DNN-based method. For the primal formulation, while the the most important signature components are those corresponding to powers of the increments, generally all components matter, and there is no apparent sparse structure. On the other hand, for the dual problem, only very few signature components actually are important, and there seems to be great potential for dimension reduction.

**Acknowledgments** 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). CB and PL acknowledge support from DFG CRC/TRR 388 “Rough Analysis, Stochastic Dynamics and Related Fields”, Project B03.

## 2 A review on Monte-Carlo methods for optimal stopping

In this section we introduce the optimal stopping problem in a fairly general framework, and recall two general simulation based procedures to derive lower and upper bounds of the optimal value. Let  $\{\Omega, \mathcal{F}, \mathbb{P}\}$  be a probability space supporting a state-process  $(X_t : 0 \leq t \leq T)$ , which we only assume to be  $\alpha$ -Hölder continuous almost surely,  $\alpha \in (0, 1)$ . Moreover, we consider an  $(\mathcal{F}_t^X)$ -adapted, continuous process  $(Z_t : 0 \leq t \leq T)$ , the cash-flow process, which for technical reasons we assume to fulfill  $\sup_{0 \leq t \leq T} |Z_t|^2 \in L^1$ . Finally, the optimal stopping problem reads

$$Y_0 = \sup_{\tau \in \mathcal{S}_0} \mathbb{E}[Z_\tau], \quad Y_t = \operatorname{ess\,sup}_{\tau \in \mathcal{S}_t} \mathbb{E}[Z_\tau | \mathcal{F}_t^X], \quad 0 < t \leq T, \quad (1)$$

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

### 2.1 A general Longstaff-Schwartz algorithm

Replacing the continuous time interval  $[0, T]$  by some finite grid  $0 = t_0 < t_1 < \dots < t_N = T$ , the *dynammic programming principle* [28, Page eq. (2.1.7)] for the discrete optimal stopping reads

$$Y_{t_N} = Z_{t_N}, \quad Y_{t_n} = \max(Z_{t_n}, \mathbb{E}[Y_{t_{n+1}} | \mathcal{F}_{t_n}^X]), \quad 0 \leq n \leq N - 1, \quad (2)$$

and  $\tau^* = \inf\{t_n : Z_{t_n} \geq \mathbb{E}[Y_{t_{n+1}} | \mathcal{F}_{t_n}^X]\}$  is optimal. Motivated by the famous Longstaff and Schwartz algorithm [25], the first optimal stopping time can be obtained as  $\tau_0$  of the following recursion

$$\tau_N = t_N, \quad \tau_n = t_n \mathbf{1}_{\{Z_{t_n} \geq \mathbb{E}[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}^X]\}} + \tau_{n+1} \mathbf{1}_{\{Z_{t_n} < \mathbb{E}[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}^X]\}}, \quad 0 \leq n \leq N - 1. \quad (3)$$

The remaining question is, how to compute the *continuation values*  $\mathbb{E}[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}^X]$  for possibly non-Markovian state-processes  $X$ , which will be the main topic of the forthcoming Section 3. For now, assume we are given a suitable family of basis-functions  $(\psi_k)$ , such that  $\mathbb{E}[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}^X] \approx \sum_k \alpha_k^n \psi_k(X|_{[0, t_n]})$  with some coefficients  $(\alpha_k^n)$ . A general version of the Longstaff-Schwartz algorithm goes as follows:

- 1 Draw  $i = 1, \dots, M$  samples  $X^{(i)}$  and  $Z^{(i)}$  and initialize  $\tau_N^{(i)} \equiv t_N$ .
- 2 Then, for  $n = N - 1, \dots, 1$

- solve the minimization problem

$$\alpha^{*,n} := \operatorname{argmin}_{\alpha} \frac{1}{M} \sum_{i=1}^M \left( Z_{\tau_{n+1}}^{(i)} - \sum_k \alpha_k \psi_k^n(X|_{[0,t_n]}^{(i)}) \right)^2$$

- Set  $\tau_n^{(i)} = t_n$  if  $Z_{t_n}^{(i)} \geq \sum_k \alpha_k^{*,n} \psi_k^n(X|_{[0,t_n]}^{(i)})$ , and  $\tau_n^{(i)} = \tau_{n+1}^{(i)}$  otherwise.

- 3 Finally, draw  $i = 1, \dots, M'$  independent sample-paths  $\tilde{X}^{(i)}$  and  $\tilde{Z}^{(i)}$ , and compute the stopping times  $\tilde{\tau}^{(i)} = \inf \left\{ t_n : \tilde{Z}_{t_n}^{(i)} \geq \sum_k \alpha_k^{*,n} \psi_k^n(\tilde{X}|_{[0,t_n]}^{(i)}) \right\}$  and the lower-biased estimator

$$\tilde{y}_0 = \max \left( Z_{t_0}, \frac{1}{M'} \sum_{i=1}^{M'} \tilde{Z}_{\tilde{\tau}^{(i)}}^{(i)} \right)$$

Notice that the resimulation in the final step ensures, that the  $\tilde{\tau}^{(i)}$  are indeed stopping times, since the computed coefficients  $\alpha$  are now independent of the samples, and thus  $\max(Z_{t_0}, \mathbb{E}[Z_{\tilde{\tau}_1}])$  is a true lower-bound. The Monte-Carlo approximation  $\tilde{y}_0$  is therefore lower-biased and strictly speaking not a true lower-bound, due to the Monte-Carlo error with respect to  $M'$ . Nevertheless, this error is typically chosen to be very small, and we will refer to  $\tilde{y}_0$  as lower-bound anyways.

## 2.2 A general dual procedure

While the Longstaff-Schwartz procedure described in the last section provides us with lower-bounds for the optimal value, it is desirable to additionally have an algorithm producing upper-bounds. To this end, it is useful to consider the (discretized) dual formulation of (1) due to [31]

$$\inf_{M \in \mathcal{M}^2} \mathbb{E} \left[ \max_{0 \leq k \leq N} (Z_{t_k} - M_{t_k}) \right], \quad (4)$$

where the minimization is over discrete  $L^2$ -martingales with respect to the filtration  $\mathcal{G}_n = \mathcal{F}_{t_n}^X$ ,  $n = 0, \dots, N$ . By the Martingale Representation Theorem [20, Theorem 4.5], if the underlying filtration  $(\mathcal{F}_t^X)$  is generated by a Brownian motion  $W$ , such martingales are of the form  $M_{t_k}^\alpha = \int_0^{t_k} \alpha_s dW_s$  for some  $(\mathcal{F}_t^X)$ -progressive process  $\alpha$ . On the other hand, due to the progressive measurability, the integrands are of the form  $\alpha_t = f(X|_{[0,t]})$  for some (deterministic) function  $f$ . Assume again that we are given some family of basis-functions  $(\psi_k)$  such that  $f(X|_{[0,t]}) \approx \sum_k \beta_k \psi_k(X|_{[0,t]})$  for some coefficients  $\beta$ . If the family  $(\psi_k)$  is *rich* enough, it is reasonable to parametrize the space  $\mathcal{M}^2$  by the span of basis-martingales  $\{M_t^k = \int_0^t \psi_k(X|_{[0,s]}) dW_s : k \geq 0\}$ , which leads to the following dual algorithm:

- 1 Draw  $i = 1, \dots, M$  sample-paths of  $Z^{(i)}, X^{(i)}, W^{(i)}$ , and compute the basis martingales  $M^{k,(i)}$  (e.g. using an Euler-scheme)
- 2 Solve the minimization problem (see, e.g., [13, 8])

$$\beta^* = \operatorname{argmin}_{\beta} \frac{1}{M} \sum_{i=1}^M \max_{0 \leq k \leq N} \left( Z_{t_n}^{(i)} - \sum_k \beta_k M_{t_n}^{k,(i)} \right)$$

- 3 Finally, draw  $i = 1, \dots, M'$  independent sample-paths  $\tilde{Z}^{(i)}, \tilde{X}^{(i)}, \tilde{W}^{(i)}$  and compute the martingales  $\tilde{M}^{k,(i)}$  and the upper-biased estimator

$$\tilde{y}_0 = \frac{1}{M'} \sum_{i=1}^M \max_{0 \leq n \leq N} \left( \tilde{Z}_{t_n}^{(i)} - \sum_k \beta_k^* \tilde{M}_{t_n}^{k,(i)} \right)$$

Similar as in the primal case, the independent resimulations ensure that we get true martingales  $\tilde{M}^* := \sum_k \beta_k^* \tilde{M}^k$ , since  $\beta^*$  is independent of the samples, and therefore  $\mathbb{E}[\max_k (Z_{t_k} - \tilde{M}_{t_k}^*)]$  is a true upper-bound. By the same abuse of terminology as in the primal case, we call the Monte-Carlo approximation  $\tilde{y}_0$  upper-bound, keeping in mind that it is strictly speaking only upper-biased due to the Monte-Carlo error with respect to  $M'$ .

### 3 Signature stopping policies and dual martingales

To numerically solve the two algorithms presented in the last section, we require to learn functionals  $f : \mathcal{X} \rightarrow \mathbb{R}$ , where  $\mathcal{X}$  is an infinite-dimensional path space. As motivated in the introduction, we will use the *path-signature*, resp. the corresponding *signature-kernel*, to tackle this problem. For the rough path theoretical details in this section, we refer to the excellent books [17, 16], but see also [10] with a focus on machine learning.

Let  $(x(t) : 0 \leq t \leq T)$  be an  $\mathbb{R}^d$ -valued,  $\alpha$ -Hölder continuous path,  $\alpha \in (0, 1]$ . The *rough path signature* of  $x$  is given by the path  $\mathbb{x}_{s,t}^{<\infty} : \Delta_{[0,T]} \rightarrow T((\mathbb{R}^d))$ , defined (at least formally) as the collection of iterated integrals

$$\mathbb{x}_{s,t}^{<\infty} := \sum_{n=0}^{\infty} \int_{s < t_1 < \dots < t_n < t} dx(t_1) \otimes \dots \otimes dx(t_n), \quad 0 \leq s \leq t \leq T,$$

taking values in the *extended tensor algebra*

$$T((\mathbb{R}^d)) := \prod_{n=0}^{\infty} (\mathbb{R}^d)^{\otimes n},$$

an algebra under a (naturally defined) non-commutative product  $\otimes$ . If  $\alpha = 1$ , the meaning of the iterated integrals is in the sense of Riemann-Stieltjes, see [17, Chapter 7.2], and less obvious becomes the case  $\alpha < 1$ , since  $dx$  does not have a meaning a-priori. However, Lyons' theory of rough paths [26] allows us to give the latter a meaning, but higher-order information of the path is required. More precisely, we have to lift  $x$  to a so-called  $\mathbf{x} = (\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}) \in \prod_{n=0}^N (\mathbb{R}^d)^{\otimes n}$ , where  $N = \lfloor \frac{1}{\alpha} \rfloor$  and  $\mathbf{x}^{(1)} = x$ , see [17, Chapter 9]. Having such a rough path lift  $\mathbf{x}$  at hand, the notion of *rough integration* allows to make sense of integrating against  $d\mathbf{x}$ , and a signature lift of  $\mathbf{x}$  can be defined, sharing all the properties of the path-signature for smooth paths, which is due to Lyons' extension theorem [26, Theorem 3.7].

**Remark 3.1** *The authors of [9] show that the map  $x \mapsto \mathbb{x}^{<\infty}$  is injective up to a certain equivalence class on path spaces, which can be eliminated when adding a monotone component to the path, e.g. the time-augmentation  $\hat{x}(t) = (t, x(t))$ . Moreover, by applying a so-called tensor-normalization  $\lambda : T((\mathbb{R}^d)) \rightarrow T((\mathbb{R}^d))$  to the signature, the authors of [12] introduce the notion of robust signatures  $\lambda(\mathbb{x}^{<\infty})$ , which can be seen as a bounded version of the signature, sharing all of its structural properties. For the theoretical results in this section, we always assume that our underlying process*

$X$  is already augmented by some monotone function, and this pair can be lifted to a geometric rough path  $\mathbf{X}$ , see [17, Definition 9.16], and we denote by  $\mathbb{X}^{<\infty}$  the unique robust rough path signature for some fixed normalization  $\lambda$ .

### 3.1 Linear signature learning

One of the most important algebraic properties of the signature is the following: For any two linear functionals on the tensor-algebra, that is  $\ell_1, \ell_2 \in T((\mathbb{R}^d)^*)$ , we can find  $\ell_3$  (in fact, a closed-form expression can be given), such that  $\langle \ell_1, \mathbb{X}_{0,T}^{<\infty} \rangle \langle \ell_2, \mathbb{X}_{0,T}^{<\infty} \rangle = \langle \ell_3, \mathbb{X}_{0,T}^{<\infty} \rangle$ , for all signatures  $\mathbb{X}_{0,T}^{<\infty}$ . In words, the space of linear functionals of the signature forms an algebra, which is point-seperating (since we assume to have a unique signature, see Remark 3.1). An application of a general Stone-Weierstrass result, see [12, Section 2.1], tells us that the set of linear functional of the robust signature is dense in the set of bounded continuous functions on the corresponding path space. This fact can be used to deduce a global  $L^p$ -approximation for linear functionals of the signature, see [6, Theorem 2.8], which then yields the following convergence results. We will always assume to be in the framework described in the beginning of Section [6, Section 3.1], where certain assumptions on the path space, as well as on the payoff process are made precise. The following proposition summarizes the results [6, Proposition 3.3 and 3.8], to which we refer for more precise statements and detailed proofs.

**Proposition 3.2** *Using the same notation from Section 2, we consider  $i = 1, \dots, M$  independent sample-paths of  $X, Z, W$ , and we define the martingales  $M_t^\ell = \int_0^t \langle \ell, \mathbb{X}_{0,s}^{\leq K} \rangle dW_s$  for some truncation level  $K \in \mathbb{N}$ . We define the estimators*

$$\begin{aligned} \ell_P^{*,n} &= \operatorname{argmin}_{\ell, |\ell| \leq K} \frac{1}{M} \sum_{i=1}^M \left( Z_{\tau_{n+1}^{(i)}}^{(i)} - \langle \ell, \mathbb{X}_{0,t_n}^{\leq K, (i)} \rangle \right)^2, \\ \ell_D^* &= \operatorname{argmin}_{\ell, |\ell| \leq K} \frac{1}{M} \sum_{i=1}^M \max_{0 \leq n \leq N} \left( Z_{t_n}^{(i)} - M_{t_n}^{\ell, (i)} \right). \end{aligned} \quad (5)$$

Then, for new samples  $\tilde{X}, \tilde{Z}, \tilde{W}$ , independent of  $X, Z, W$ , define the  $(\mathcal{F}_t^X)$ -stopping times  $\tilde{\tau}^{(i)} = \inf \left\{ t_n : \tilde{Z}_{t_n}^{(i)} \geq \langle \ell_P^{*,n}, \mathbb{X}_{0,t_n}^{\leq K, (i)} \rangle \right\}$  and martingales  $\tilde{M}^{\ell_D^*, (i)}$ . Then

$$\max \left( \tilde{Z}_{t_0}, \frac{1}{M} \sum_{i=1}^M \tilde{Z}_{\tilde{\tau}^{(i)}}^{(i)} \right) \nearrow Y_0, \quad \frac{1}{M} \sum_{i=1}^M \max_{0 \leq n \leq N} \left( \tilde{Z}_{t_n}^{(i)} - \tilde{M}_{t_n}^{\ell_D^*, (i)} \right) \searrow Y_0, \quad (6)$$

as  $M, K \rightarrow \infty$ , where the convergence with respect to  $M$  is almost sure convergence.

While the last result ensures converges for the primal and dual procedures, is does not come with quantitative statements about convergence rates. Additionally, these algorithms can become computationally expensive, as the size of the signature grows exponentially with  $K$ . In [6] several numerical experiments were performed for these two algorithms in non-Markovian frameworks, in particular for models driven by fractional Brownian motion with small Hurst parameters. The accuracy of the method can be measured by the duality gap between lower and upper bounds, and in some examples these gaps were observed to be quite significant even for high truncation levels. Two important sources for this gap can be described as follows: First, especially in very rough regimes, the considered signature levels might not suffice to capture the relevant past of the non-Markovian processes, so that information is lost when truncating the signature. Secondly, the functionals on the path spaces we try to

learn are highly non-linear, so that more sophisticated learning technologies are required to improve the performance. The goal of the next sections is to introduce non-linear extensions of the primal and dual algorithm, based on deep-, resp. kernel-learning methodologies, with the objective of improving the duality gap.

### 3.2 Deep signature learning

Applying *Deep Neural Networks* (DNNs) on top of the signature was considered several times in the literature, see for instance [21], where even more generally, the signature can act as a layer in a network. Here, we simply consider the truncated signature as the input for a DNN of the form

$$\theta = \beta \circ \sigma \circ A_1 \circ \sigma \cdots \circ A_I,$$

where  $\beta, A_0, \dots, A_I$  are affine,  $\beta : \mathbb{R}^m \rightarrow \mathbb{R}$ ,  $A_I : T(\mathbb{R}^d) \rightarrow \mathbb{R}^m$ ,  $A_{I-i} : \mathbb{R}^m \rightarrow \mathbb{R}^m$ , for  $i = 1, \dots, I-1$ , and  $\sigma$  is an activation function. We call  $\theta$  a *signature DNN* with  $I$  hidden layers, each of which having  $m$  neurons. The signature map  $\mathbf{x}|_{[0,t]} \mapsto \mathbb{x}_{0,t}^{<\infty}$  represents the input layer, and the first hidden layer acting on the signature, can be seen as a vector of linear functionals  $A_I = (\ell_1, \dots, \ell_m)$ , and  $A_I(\mathbb{x}^{<\infty}) = (\langle \ell_1, \mathbb{x}^{<\infty} \rangle, \dots, \langle \ell_m, \mathbb{x}^{<\infty} \rangle)^\top \in \mathbb{R}^m$ . The remaining hidden layers are all of the form  $A_j(x) = A_j x + b_j$ , where  $A_j \in \mathbb{R}^{m \times m}$ ,  $b_j \in \mathbb{R}^m$ , and the output layer is given by  $\beta \in \mathbb{R}^m$ . We denote by  $\text{DNN}_{\text{sig}}^{\sigma, I}$  the set of all such signature DNNs  $\theta$  with  $I$  hidden layers, and activation function  $\sigma$ .

It is well known, that already  $I = 1$  hidden layer DNNs are universal, see, e.g. [24]. For this reason, and to ease the notation for the theoretical results, in the remainder of this section we fix  $I = 1$ . Nevertheless, all the results trivially extend to multiple hidden layers  $I > 1$ , and in the numerical experiments we mostly choose more than one hidden layer. Setting  $\text{DNN}_{\text{sig}}^\sigma := \text{DNN}_{\text{sig}}^{\sigma, 1}$ , we can explicitly write the set as

$$\text{DNN}_{\text{sig}}^\sigma = \left\{ \mathbf{x}|_{[0,t]} \mapsto \sum_{j=1}^m \beta_j \sigma \left( \langle \ell_j, \mathbb{x}_{0,t}^{<\infty} \rangle \right) : (\beta, \ell) \in \mathbb{R}^m \times \mathcal{W}^m, m \geq 1 \right\}, \quad (7)$$

where  $\mathcal{W}$  denotes the set of all linear functionals, and  $\mathbf{x}|_{[0,t]}$  are *stopped  $\alpha$ -Hölder rough paths*, see [6, Section 2] for details. For fixed  $(m, k) \in \mathbb{N}^2$  and any pair  $(\beta, \ell) \in \mathbb{R}^m \times (\mathcal{W}^{\leq k})^m$ , where  $\mathcal{W}^{\leq k}$  are linear functionals on the  $k$ -step signature  $\mathbb{x}^{\leq k}$ , denote by  $\theta^{(\beta, \ell)}$  the corresponding functional in the set  $\text{DNN}_{\text{sig}}^\sigma$ . Revisiting the Longstaff & Schwartz algorithm presented in Section 2.1, it is tempting to define a sequence of *deep stopping times*  $(\tau_n^{(m, k)} : 1 \leq n \leq N)$ , recursively defined similar to (3), but with conditional expectations replaced by  $\theta^{(\beta_P^{*, n}, \ell_P^{*, n})}(\mathbf{X}|_{[0, t_n]})$ , where for  $0 \leq n \leq N$

$$(\beta_P^{*, n}, \ell_P^{*, n}) = \underset{(\beta, \ell) \in \mathbb{R}^m \times (\mathcal{W}^{\leq k})^m}{\text{argmin}} \mathbb{E} \left[ \left( Z_{\tau_{n+1}^{(k, m)}} - \theta^{(\beta, \ell)}(\mathbf{X}|_{[0, t_n]}) \right)^2 \right]. \quad (8)$$

In words, we recursively learn the continuation values as neural networks of truncated signatures. Similarly, we approximate the Doob-martingale from Section 2.2 by *deep martingales*  $M_t^{(k, m)} = \int_0^t \theta^{(\beta_D^*, \ell_D^*)}(\mathbf{X}|_{[0, s]}) dW_s$ , where

$$(\beta_D^*, \ell_D^*) = \underset{(\beta, \ell) \in \mathbb{R}^m \times (\mathcal{W}^{\leq k})^m}{\text{argmin}} \mathbb{E} \left[ \max_{0 \leq n \leq N} \left( Z_{t_n} - \int_0^{t_n} \theta^{(\beta, \ell)}(\mathbf{X}|_{[0, s]}) dW_s \right) \right]. \quad (9)$$

Thanks to the universality of both signatures and DNNs, the following result should not come as a surprise. Similar as in the last section, we adopt the assumptions described in the beginning of [6, Section 3.1].

**Proposition 3.3** *Assuming that the filtration is Brownian, that is  $\mathcal{F}^X = \mathcal{F}^W$ , and the activation function  $\sigma$  is non-polynomial, we have*

$$Y_0^N = \lim_{k,m \rightarrow \infty} \max \left( Z_{t_0}, \mathbb{E}[Z_{\tau_1^{(k,m)}}] \right) = \lim_{k,m \rightarrow \infty} \mathbb{E} \left[ \max_{0 \leq n \leq N} (Z_{t_n} - M_{t_n}^{(k,m)}) \right].$$

An outline of the proof can be found in Appendix A. Similar to the linear case, in practice we solve the sample average approximation of (8)-(9), that is

$$\begin{aligned} (\beta_P^{*,n}, \ell_P^{*,n}) &= \operatorname{argmin}_{(\beta, \ell) \in \mathbb{R}^m \times (\mathcal{W}^{\leq k})^m} \frac{1}{M} \sum_{i=1}^M \left( Z_{\tau_{n+1}^{(i)}}^{(i)} - \theta^{(\beta, \ell)}(\mathbf{X}^{(i)}|_{[0, t_n]}) \right)^2, \\ (\beta_D^*, \ell_D^*) &= \operatorname{argmin}_{(\beta, \ell) \in \mathbb{R}^m \times (\mathcal{W}^{\leq k})^m} \frac{1}{M} \sum_{i=1}^M \max_{0 \leq n \leq N} \left( Z_{t_n}^{(i)} - M_{t_n}^{(\beta, \ell), (i)} \right), \end{aligned} \quad (10)$$

for i.i.d sample-paths  $Z^{(i)}, \mathbf{X}^{(i)}, W^{(i)}$ . The latter (non-convex) optimization problems can then be solved using stochastic gradient descent. More details about this and the DNN architecture will be discussed in Section 4. Let us conclude this section with a remark.

**Remark 3.4** *To address the complexity issue coming from the size of the signature, it can be helpful to consider the so-called log-signature. The latter is defined by  $\mathbb{L}_{s,t}^{\leq \infty} := \log^{\otimes}(\mathbb{X}_{s,t}^{\leq \infty})$ , where  $\log^{\otimes}$  is the bijection*

$$\log^{\otimes}(1 + \mathbf{x}) := \sum_{k \geq 0} \frac{(-1)^{k+1}}{k} \mathbf{x}^{\otimes k} \in T((\mathbb{R}^d)), \quad \mathbf{x} \in T((\mathbb{R}^d)), \mathbf{x}^{(0)} = 0.$$

*It can be shown that the dimension of the truncated log-signature  $\mathbb{L}^{\leq K}$ , grows much slower than the one of  $\mathbb{X}^{\leq K}$ , see, e.g., [29], and denoting by  $\exp^{\otimes}$  the inverse of  $\log^{\otimes}$ , we have  $\mathbb{X}^{\leq K} = \exp^{\otimes}(\mathbb{L}^{\leq K})$ . Although the log-signature itself is not longer universal, Proposition 3.3 remains through for log-signature DNNs, that is replacing the signature in (7) by its log-signature transform. In Remark A.1, after the proof of Proposition 3.3, we briefly explain why this is true and how to modify the proof.*

### 3.3 Signature-kernel learning

Let us start by recalling some notions from general RKHS theory, see, e.g. [34]. Given a feature map  $\phi : \mathcal{X} \rightarrow \mathcal{H}$ , where  $\mathcal{H}$  is a Hilbert space (the so-called *feature space*), one can define the associated kernel

$$k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}, \quad k(x, y) := \langle \phi(x), \phi(y) \rangle_{\mathcal{H}}. \quad (11)$$

If the kernel is positive-definite, there exists a unique RKHS with reproducing kernel  $k$ , see [34, Theorem 4.21], in the sense that  $k$  reproduces elements  $f \in \mathcal{H}$ , that is  $f(x) = \langle k(\cdot, x), f \rangle_{\mathcal{H}}$ , and it generates the Hilbert space,  $\mathcal{H} = \overline{\operatorname{span}\{k(\cdot, x) : x \in \mathcal{X}\}}$ . In our case, the feature map  $\phi$ , maps a path  $x \in \mathcal{X}$  to its path-signature  $\mathbb{x}^{\leq \infty}$ , and therefore signature kernel is naturally defined by

$$k_{s,t}(x, y) := \langle \mathbb{x}_{0,s}^{\leq \infty}, \mathbb{y}_{0,t}^{\leq \infty} \rangle, \quad 0 \leq s, t \leq T$$

where  $\langle \cdot, \cdot \rangle$  is the natural extension of the inner products  $(\mathbb{R}^d)^{\otimes k}$  to the tensor algebra  $T((\mathbb{R}^d))$ , see [10, Definition 2.1.1]. It has been shown in [32], that the signature-kernel solves a Goursat-type PDE with respect to the time variables. This *kernel trick* allows us to evaluate the signature-kernel,

representing the whole signature, by numerically solving a second-order PDE. An important extension, designed for rougher inputs, was discussed in the recent work [23].

An important observation is that universality of the signature feature map is equivalent to the universality of the signature-kernel, see for instance [12, Section 6], but also [10, Chapter 2.1]. Returning to optimal stopping, this motivates us to seek for functionals in the primal and dual problems, solving the regularized minimizing problems on the RKHS

$$f^\lambda = \operatorname{argmin}_{f \in \mathcal{H}} \mathcal{L}(f) + \lambda \|f\|_{\mathcal{H}}^2, \quad \lambda > 0, \quad (12)$$

where the loss  $\mathcal{L}$  is either the primal or the dual loss function. More precisely, in the primal case, at each exercise date  $n$ , the loss is simply the mean-square error  $\mathcal{L}_n(f) = \|Z_{t_n} - f(X|_{[0,t_n]})\|_2^2$ . In the dual case, it is given by  $\mathcal{L}(f) = \|\max_{1 \leq k \leq N} (Z_{t_k} - M_{t_k}^f)\|_1$ , where  $M_t^f = \int_0^t f(X|_{[0,s]}) dW_s$ . Then, replacing the losses by their sampled version, the general representer theorem [33, Theorem 1] reveals that the minimizers are of the form  $f^\alpha(x) = \sum_{i=1}^M \alpha_i k(X^{(i)}, x)$ . Thus, let us define the minimizers

$$\alpha_P^{*,n} := \operatorname{argmin}_{\alpha \in \mathbb{R}^M} \frac{1}{M} \sum_{i=1}^M \left( Z_{\tau_{n+1}}^{(i)} - \sum_{j=1}^M \alpha_j k_{t_n, t_n}(X^{(i)}, X^{(j)}) \right)^2 + \lambda \|f^\alpha\|_{\mathcal{H}}^2, \quad \lambda > 0, \quad (13)$$

for all exercise date  $n = 1, \dots, N - 1$ . Similarly,

$$\alpha_D^* := \operatorname{argmin}_{\alpha \in \mathbb{R}^M} \frac{1}{M} \sum_{i=1}^M \max_{0 \leq n \leq N} \left( Z_{t_n}^{(i)} - \sum_{j=1}^M \alpha_j M_{t_n}^{(i),(j)} \right)^2 + \lambda \|f^\alpha\|_{\mathcal{H}}^2, \quad \lambda > 0, \quad (14)$$

for the kernel-martingales  $M_t^{(i),(j)} = \left( \int_0^t k_{s,s}(X^j, X) dW_s \right)^{(i)}$ . The algorithms in Section 2 can then easily be translated to the kernel-learning framework, by replacing the minimizers in the second step accordingly.

**Remark 3.5** We expect that the convergence results in Proposition 3.2 can be obtained analogously in this kernel-learning framework, when sending  $M \rightarrow \infty$  and  $\lambda \rightarrow 0$ . For instance in the dual case, it follows by universality of the signature-kernel, that the closure of the span of the family of kernel-martingales  $\mathcal{M}^K = \left\{ \int k_{s,s}(x, X) dW_s : x \in \mathcal{X} \right\}$  corresponds to the space of  $L^2(\mathcal{F}^W)$ -martingales. For the sample average approximation and existence of minimizers, one can then argue similar as in [6, Appendix A.2.], and use the representer theorem to conclude.

**Remark 3.6** The obvious advantage of this method is that no truncation for the feature map is necessary, and therefore theoretically it does not suffer from a loss of information, see Remark 3.1. Modulo evaluation of the signature-kernel, the approximation error only depends on the regularization  $\lambda$  and the number of samples  $M$ . Moreover, at least for the mean-square loss, it seems possible to theoretically study convergence rates of such algorithms. This, however, involves a more precise understanding of the signature RKHS  $\mathcal{H}$  and integral operators therein, a problem that is outside the scope of this paper, but planned for future research.

Having mentioned the theoretical advantages of the kernel-method, evaluating the signature-kernel, which means solving the Goursat-PDE [32], becomes the main difficulty in this procedure. Due to the recursive nature of the regression problems in the Longstaff and Schwartz algorithm in Section 2.1, typically large sample size are required to ensure stability. Moreover, as we are especially interested

in paths of low regularity, a fine discretization grid is required to solve the Goursat PDE. Combining these observations with the fact, that the kernel-ridge-regressions involve computing and inverting the Gram-matrices  $(k_{t_n, t_n}(X^{(i)}, X^{(j)}))_{i,j}$ , it becomes clear that a reduction of the computational costs is required.

To this end, we propose the following approach, related to the so-called *Nyström-method* for kernel-learning [14, 37]. We randomly select  $L \ll M$  subsamples, denoting by  $I_L$  the set of those indices, and in both the primal and dual optimization problems we restrict the minimization to functions of the form  $f(\cdot) = \sum_{i \in I_L} \beta_i k(x^i, \cdot)$ . It follows that we only need to compute the  $(M \times L)$ -matrices

$$\mathbf{K}_{t_n} = (k_{t_n, t_n}(X^j, X^i) : i \in I_L, j = 1, \dots, M), \quad n = 1, \dots, N.$$

For the primal example, that is, for the kernel ridge regression, we can additionally observe that the explicit solutions is given by

$$\alpha_P^{*,n} = (\mathbf{K}_{t_n}^T \mathbf{K}_{t_n} + M\lambda \mathbf{R}_{t_n})^{-1} \mathbf{K}_{t_n}^T Z_{\tau_{n+1}}$$

where  $\mathbf{R}_{t_n} = (k_{t_n, t_n}(x^i, x^j) : i, j \in I_L)$ . Further details about the implementation will be presented in the next section.

## 4 Pricing American options under rough volatility

In this section we test the signature-based procedures for the problem of pricing American options in rough volatility models. In such models, the asset-price dynamics is given by

$$S_0 = s_0, \quad dS_t = rS_t dt + S_t v_t \left( \rho dW_r + \sqrt{1 - \rho^2} dB_t \right), \quad 0 < t \leq T, \quad (15)$$

where  $W$  and  $B$  are two independent Brownian motions, the volatility  $(v_t)_{t \in [0, T]}$  is  $(\mathcal{F}_t^W)$ -adapted and continuous,  $\rho \in [-1, 1]$  and  $r > 0$  the interest rate. We denote by  $X$  the log-price, that is  $X_t = \log(S_t)$ . Compared to classical diffusion models, we are interested in volatility process  $(v_t)$  driven by fractional Brownian motion, turning both the volatility and the price into non-Markovian processes. We will focus on the following two, arguably most popular examples of rough volatility specifications.

**Example 1** (Rough Bergomi [3]) The rough Bergomi volatility is given by

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

where  $\mathcal{E}$  denotes the stochastic exponential. In all the numerical examples, we will choose  $\eta = 1.9$  and  $\xi_0 = 0.09$ .

**Example 2** (Rough Heston [15]) The rough Heston volatility, resp. its variance  $V = v^2$ , is defined as (weak) solution to the Volterra-tye CIR equation

$$V_t = V_0 + \int_0^t (t-s)^{H-\frac{1}{2}} \lambda (\theta - V_s) ds + \int_0^t (t-s)^{H-\frac{1}{2}} \nu \sqrt{V_s} dW_s. \quad (16)$$

In all the numerical examples, we will choose  $V_0 = \theta = 0.02$ ,  $\nu = \lambda = 0.3$ .

The problem of pricing American (resp. Bermudan) options consists of solving the discrete optimal stopping problem

$$Y_0^N = \sup_{\tau \in \mathcal{S}_0^N} E[e^{-r\tau} \phi(S_\tau)], \quad (17)$$

where  $\mathcal{S}_0^N$  is the set of  $(\mathcal{F}_t^S)$ -stopping times taking values in the set of possible exercise-dates  $\{t_0, \dots, t_N\}$ , and  $\phi$  is the payoff function. We will focus on so-called *put-options*, which corresponds to the payoff function  $\phi(x) = (K - x)^+$  for a given strike price  $K$ .

## 4.1 Implementation details

The code accompanying this section can be found in [https://github.com/lucapeizzari/Optimal\\_Stopping\\_with\\_signatures](https://github.com/lucapeizzari/Optimal_Stopping_with_signatures). Before discussing numerical experiments, let us specify in detail how the procedures introduced in Section 3 are implemented. All the signatures are computed using the package *iisignature*, see [30]. For the signature kernel, we rely on the PDE solvers from the package *sigkernel* to obtain the kernel as finite-difference solution to the Goursat PDE derived in [32]. For the simulation of the rough Bergomi, we use [https://github.com/ryanmccrickerd/rough\\_bergomi](https://github.com/ryanmccrickerd/rough_bergomi) related to [27], and for the simulation of rough Heston [https://github.com/SimonBreneis/approximations\\_to\\_fractional\\_stochastic\\_volterra\\_equations](https://github.com/SimonBreneis/approximations_to_fractional_stochastic_volterra_equations) related to [2].

### Linear signature stopping

This approach was already studied in detail in [6], and we adopt the choices made there. In particular, for the primal procedure we use the signature of  $\hat{X}_t = (t, X_t)$ , and for the dual procedure the signature of  $\hat{Z}_t = (t, X_t, \phi(X_t))$ , and in both we add Laguerre polynomials  $L(X_t, v_t)$  to the set of basis functions. For more details about this choice we refer to [6, Section 4].

### Deep signature stopping

Here we make a slightly different choice for the basis compared to the linear approach, exploiting both the universality of the DNNs and the signature. First, it was observed for instance in [1], that the price  $X$  has a partial Markovian nature in  $B$ , and only depends on the past through the non-Markovian volatility process  $(v_t)_{t \in [0, T]}$ . Leaving rigorous arguments to [1], one can for instance note that  $\text{Law}(X_{t+h} | \mathcal{F}_t) = \text{Law}(X_{t+h}^{t,x} | \mathcal{F}_t^v)|_{x=X_t}$ , so that conditional expectations  $\mathbb{E}[\phi(X_{t+\Delta t}) | \mathcal{F}_t] = f(X_t, (v_s)_{s \leq t})$  for some measurable function  $f$ . To capture the relevant memory of the dynamics, in the primal case we lift the variance-augmented process  $\hat{v}_t = (\langle X \rangle_t, v_t) = (\int_0^t v_u^2 du, v_t)$  the the standard signature  $\hat{\mathbb{V}}^{<\infty}$ . In the dual case we simply lift the time-augmentation  $\hat{v}_t = (t, v_t)$ . The partial Markovianity motivates to simply apply DNNs on  $\{X_t, \hat{\mathbb{V}}_{0,t}^{\leq K}\}$ .

Let us now specify the DNN architecture: For the Longstaff and Schwartz algorithm, similar to [22], we rely on the *Leaky ReLu* activation function  $\sigma(x) = 1_{\{x \geq 0\}}x + 1_{\{x < 0\}}0.3x$  and the ADAM optimizer to fit the models at each exercise date, with a batch-size of  $b = 128$  and learning rate  $\lambda = 10^{-3}$ . Inspired by [22], we use  $e = 15$  epochs to learn the conditional expectation at the last exercise date (first regression in the algorithm), and then use the trained weights to initialize the DNN at the next exercise-date. Doing this allows to reduce the epochs to  $e = 1$  for  $1 \leq n < N - 1$ , which in turn reduces the computation time significantly. It is sufficient to consider  $I = 2$  hidden layers with each

$D_i \equiv (\dim(\hat{\mathbb{V}}^{\leq K}) + 1) + 32$  neurons. For the dual algorithm we rely on the classical *ReLU* activation function  $\sigma(x) = \max(x, 0)$ , and again using the ADAM optimizer with batch size  $b = 128$  and learning rate  $10^{-3}$  to fit the model. Compared to the primal algorithm, deeper neural networks with  $I = 6$  layers and  $D_i \equiv (\dim(\hat{\mathbb{V}}^{\leq K}) + 1) + 32$  are required, which is due to the non-linear nature of the integrand of the Doob-martingale.

### Signature-kernel stopping

For the kernel ridge regressions in the Longstaff and Schwartz algorithm based methods, we choose the kernel of the time-augmented path  $\hat{N}_t = (\langle X \rangle_t, W_t, X_t)$ . At each exercise date  $n$ , we randomly select  $L = 32$  subsamples, where the  $i$ -th sample is selected with probability  $p_i = \frac{k_{t_n, t_n}(\hat{N}^{(i)}, \hat{N}^{(i)})}{\sum_{j=1}^M k_{t_n, t_n}(\hat{N}^{(j)}, \hat{N}^{(j)})}$ , noting that the latter only requires the diagonal of the Gram matrix  $\mathbf{K}$ .

For the dual procedure, we randomly select the subsamples according to the probabilities, build on the quadratic variation of the kernel martingales  $\langle \int_0^T k_{s,s}(\hat{N}^{(i)}, \hat{N}) dW_s \rangle_T^{(i)} = \int_0^T k_{s,s}(\hat{N}^{(i)}, \hat{N}^{(i)})^2 ds$ , that is

$$p_i = \frac{\int_0^T k_{s,s}(\hat{N}^{(i)}, \hat{N}^{(i)})^2 ds}{\sum_{j=1}^M \int_0^T k_{s,s}(\hat{N}^{(j)}, \hat{N}^{(j)})^2 ds}.$$

Having  $\mathbf{K} \in \mathbb{R}^{M \times L}$  at hand, we noticed that best performance can be achieved when solving (14) using a simple version of the neural network technology before, with 1 hidden-layer and *ReLU* activation.

## 4.2 American put option prices in rough Bergomi and rough Heston

In this section we compare all the signature-based methods to price Bermudan put options. In Table 1-2, we present the rough Bergomi model with Hurst parameters  $H = 0.8$ , resp.  $H = 0.07$ , and for a range of strikes  $K \in \{70, 80, \dots, 120\}$  and  $S_0 = 100$ . We choose the contract duration of  $T = 1$  year and  $J = 12$  early exercise options and interest rate  $r = 0.05$  and correlation  $\rho = -0.9$ . In the first column we report the European price, that is, the price for no early exercise opportunity,  $\mathbf{E} = \mathbb{E}(e^{-rT}(K - S_T)^+)$ . The second columns correspond to lower-bounds obtained in [18], see also [7]. Finally, in the remaining columns, we compare the *point estimate*, which is the biased-estimator from the Longstaff and Schwartz algorithm based on the training samples, with the lower- and upper-bounds obtained from independent testing samples. While the intervals in the linear case are taken from [6], we derive the intervals for the deep signature using truncation level  $K = 4$ , discretization for the signature and martingales  $N = 600$ , and both  $M = 12^{18}$  samples for training and testing, to ensure stability of the procedure. For the signature-kernel, we use  $M = 2^{17}$  samples, with  $N = 240$  discretization points for solving the Goursat PDE. For each strike, we highlight the best lower-, resp. upper-bound, and in the last column present the relative duality gap with respect to these two values. In Table 3 we present the same considerations for the rough Heston model with  $H = 0.1$ ,  $r = 0.06$  and  $\rho = -0.7$ , a choice motivated in [2], with smaller regularization parameter. The trainings were repeated 20 times and the overall Monte-Carlo error is below 0.01. Finally, in Table 4 we compare the computational time (in seconds) of one training each, as well as the offline computation of the signature, resp. signature-kernel. In Figure 1 we present the dependence of the lower-bounds, resp. point estimates with respect to  $\lambda \in [0, 2]$ . The optimal choice is made for highest possible lower bounds (blue lines), so that the distance to the point estimate is reasonable, since larger difference

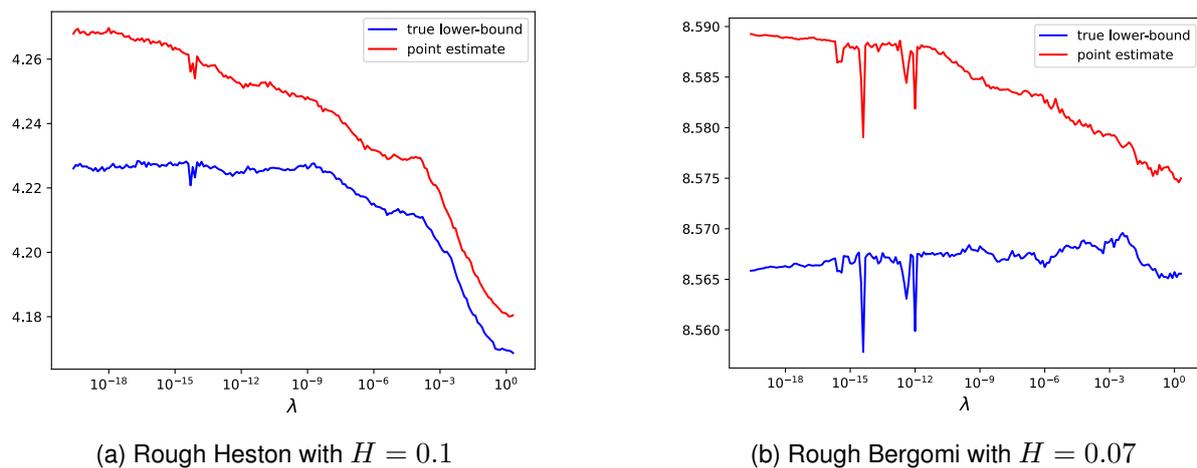


Figure 1: Primal put option prices with strike  $K = 100$  and signature-kernel procedure, with respect to the penalization parameters  $\lambda$ .

suggest overfitting. According to this remark, we choose  $\lambda = 10^{-3}$  for rough Bergomi, and  $\lambda = 10^{-8}$  in the rough Heston model.

First and most visibly, we observe that the the deep signature method delivers significantly better upper-bounds than both other methods, with reasonable training and offline costs, see Table 4. Moreover, the deep method also mostly produces the smallest duality gap, that is the smallest distance between lower- and upper-bound. However, one can also observe that for both the point estimates and lower-bounds, the linear and kernel methods, which are based on conventional convex optimization problems, slightly outperform the deep signature method in general. It should be noted that, however, this comes at the price of offline computational costs, see Table 4, where we recall that the DNNs allow to only lift the augmented volatility process. As discussed in [6], for the linear procedure it is necessary to lift the three-dimensional path  $(t, X_t, \phi(X_t))$  and add polynomials of the state, which explains the increased computational time in the linear case. While the signature-kernel method improves the linear approach, the complexity of deriving the Gram-matrix, i.e. solving the Goursat-PDE, see Table 4, leaves the signature-kernel procedure as the most expensive approach to solve the optimal stopping problem. To handle large data sets, it is therefore necessary to solve the complexity issue for the kernel, and a potential direction could be the *random Fourier signature features* by [36], but this is outside the scope of this paper.

In summary, for our goal of achieving minimal duality gaps for arbitrary large sample sizes, the best possible results can be achieved by combining a Longstaff & Schwartz algorithm based on the linear signature or the signature-kernel, exploiting the simplicity of the training procedure, together with the deep-signature approach for the upper-bounds, to capture the non-linearity of the Doob martingale integrand. If one has to choose one method, we recommend the deep-signature method, simply as it shows the smallest duality gaps and overall computational times.

On the other hand, it is important to note that we observe the primal signature-kernel method to be more stable compared to deep-signatures, with respect to smaller data sets. This unsurprising advantage of the kernel-method can become important when one works with real-word rather than synthetic data, where the number of samples cannot be arbitrary large, in which case the signature-kernel might be favorable. This is illustrated in Figure 2 (a) for the rough Bergomi model with  $H = 0.07$ , where we can observe more severe instability for the neural network training for smaller  $M$ , both in the sense of small lower-bounds and high training variance. The Monte-Carlo errors reflect this

Strike	E	[18]	Linear [6]		Deep signature		Signature kernel		Best Gap
			Point est.	C.I.	Point est.	C.I.	Point est.	C.I.	
70	0.78	1.84	1.847	[1.83, 1.90]	1.831	[1.82, 1.85]	1.835	[1.83, 1.89]	1.08 %
80	1.55	3.10	3.101	[3.08, 3.19]	3.086	[3.07, 3.12]	3.080	[3.08, 3.16]	1.28 %
90	3.11	5.08	5.086	[5.07, 5.17]	5.058	[5.04, 5.08]	5.061	[5.06, 5.15]	0.78 %
100	6.10	8.19	8.188	[8.15, 8.27]	8.132	[8.11, 8.17]	8.159	[8.15, 8.26]	0.24 %
110	11.41	13.00	12.991	[12.97, 13.09]	12.922	[12.87, 12.98]	12.944	[12.94, 13.03]	0.07 %
120	18.65	20.28	20.219	[20.21, 20.51]	20.161	[20.14, 20.25]	20.162	[20.16, 20.27]	0.19 %

Table 1: Comparison of Put option prices in rough Bergomi with  $H = 0.8$ 

Strike	E	[18]	Linear [6]		Deep signature		Signature kernel		Best Gap
			Point est.	Interval	Point est.	Interval	Point est.	Interval	
70	0.88	1.88	1.929	[1.92, 1.99]	1.921	[1.91, 1.95]	1.926	[1.92, 2.01]	1.53 %
80	1.67	3.25	3.289	[3.27, 3.37]	3.281	[3.26, 3.31]	3.286	[3.27, 3.36]	1.20 %
90	3.11	5.34	5.394	[5.37, 5.50]	5.383	[5.35, 5.44]	5.397	[5.37, 5.53]	1.28 %
100	5.83	8.53	8.586	[8.57, 8.77]	8.555	[8.52, 8.66]	8.589	[8.56, 8.75]	1.03 %
110	10.94	13.28	13.314	[13.29, 13.59]	13.281	[13.21, 13.45]	13.326	[13.27, 13.46]	1.18 %
120	18.37	20.20	20.267	[20.24, 20.66]	20.163	[20.14, 20.63]	20.276	[20.24, 20.44]	0.97 %

Table 2: Comparison of Put option prices in rough Bergomi with  $H = 0.07$ .

Strike	E	Linear [6]		Deep signature		Signature kernel		Best Gap
		Point est.	Interval	Point est.	Interval	Point est.	Interval	
70	0.16	0.438	[0.42, 0.53]	0.426	[0.42, 0.44]	0.430	[0.43, 0.46]	2.27 %
80	0.42	0.966	[0.94, 1.11]	0.961	[0.95, 0.97]	0.960	[0.96, 1.01]	1.03 %
90	1.02	2.034	[1.99, 2.28]	2.031	[2.02, 2.07]	2.027	[2.02, 2.11]	2.42 %
100	2.66	4.245	[4.16, 4.66]	4.248	[4.24, 4.36]	4.219	[4.21, 4.44]	2.75 %
110	8.05	9.715	[9.63, 10.52]	9.660	[9.66, 10.42]	9.693	[9.68, 10.14]	4.53 %
120	16.68	19.500	[19.49, 20.04]	19.487	[19.48, 20.00]	19.502	[19.50, 19.56]	0.30 %

Table 3: Comparison of Put option prices in rough Heston with  $H = 0.1$ .

Method	Training primal	Training dual	Offline costs
Linear	1.73	610.10	508.32
Deep	100.23	520.34	50.92
Kernel	1.69	390.98	5564.53

Table 4: Computational times (in seconds) for training and evaluation, and offline computation times (in seconds) for signatures and kernels

variance, which is obtained after independently repeating the training 20 times for each sample size. For completeness, we show a similar plot in the dual case in Figure 2 (b), where it is confirmed again that the deep-signature method heavily outperforms the signature-kernel approach, also in every data-size regime. The variance in the dual procedure is much smaller, as we only solve one optimization problem in the training, rather than at each exercise date.

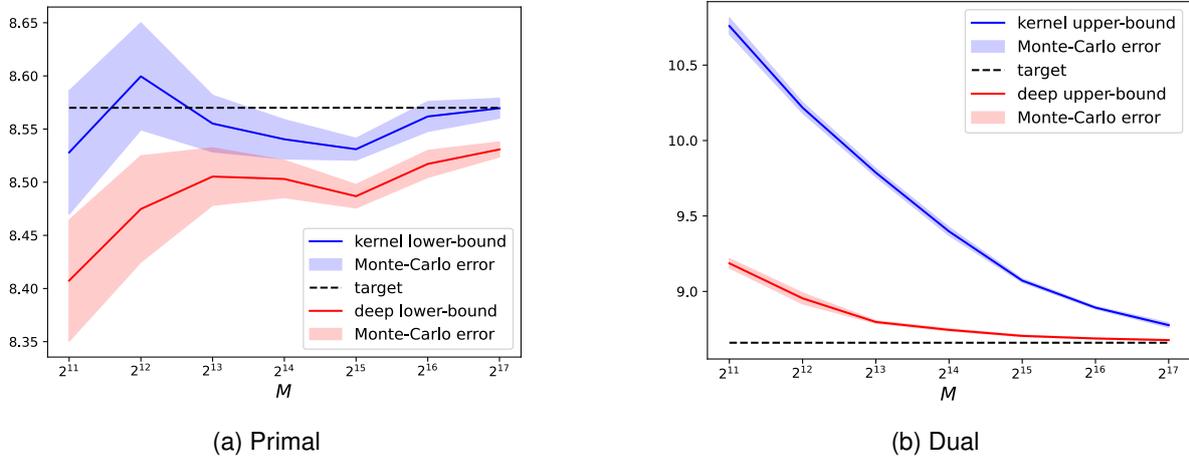


Figure 2: Put option prices for  $H = 0.07$  and strike  $K = 100$ , with respect to training sample size  $M$ .

In Figure 3 we compare the options prices in the rough Bergomi model, with respect to the correlation parameter  $\rho \in [-1, 1]$ , for the two (very different regimes) regimes  $H = 0.07$  and  $H = 0.8$ . The blue region reflects the pricing interval derived from the deep-signature approach, and once again we present the point-estimates of all methods. In both cases we can observe, that the pricing method perform worse – in the sense of duality gap – when we come closer to the no-correlation case  $\rho = 0$ , while the intervals get tighter in more correlated regimes.

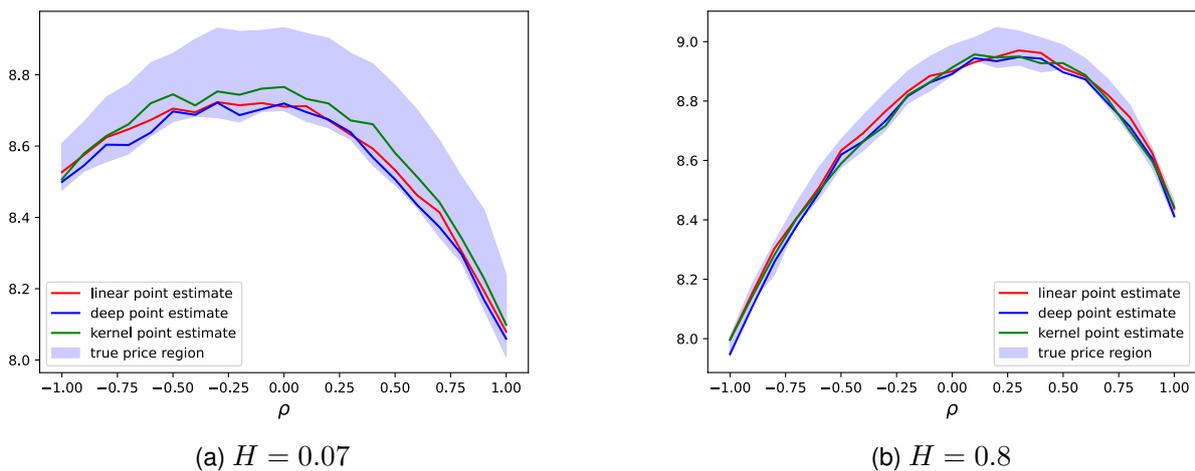


Figure 3: Point-estimates and pricing intervals for put option with  $K = 100$  in the rough Bergomi model, with respect to the correlation parameters  $\rho \in [-1, 1]$ .

In Figure 4 we compare the duality gap in the deep-signature method, with respect to the number of

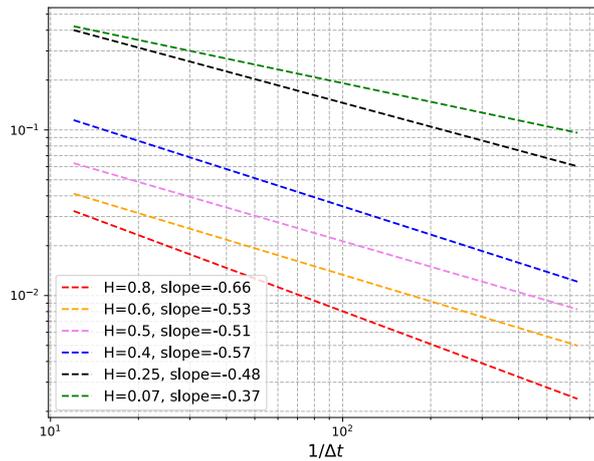


Figure 4: Linearly fitted duality gaps with respect to time-discretization and different Hurst regimes in the rough Bergomi model.

discretization points and different Hurst regimes. More precisely, let us define the relative duality gap

$$\epsilon = \epsilon(N, H) = \frac{y_0^D(N, H) - y_0^P(N, H)}{y_0^D(N, H)},$$

where  $y_0^D(N, H)$  (resp.  $y_0^P(N, H)$ ) is the upper (resp. lower) bound obtained from the deep-signature method in the rough Bergomi model with Hurst parameter  $H \in (0, 1)$  and  $N \in \mathbb{N}$ . All the other parameters are chosen the same as in the Tables 1-2. We can observe that in general higher Hurst parameters show smaller duality gaps and also faster convergence with respect to  $N$ . Notice that, apart from optimization and Monte-Carlo errors, the two main sources for the gap here are the discretization error from approximating iterated and stochastic integrals, and the non-Markovianity for  $H \neq 1/2$  regimes. When  $H$  is close to  $1/2$ , the convergence rate seems to be around  $H$ , and since this regime is "almost Markovian" (and Markovian for  $H = 1/2$ ), this can be interpreted as the strong error occurring when approximating the stochastic integrals  $\int f(v_t, X_t) dW_t$  with an Euler-scheme and rough integrand. In general, however, it is not possible to separate these two error sources, so that it becomes more difficult to interpret the observed convergence rate. Finally, in Figure 5 we compare the impor-

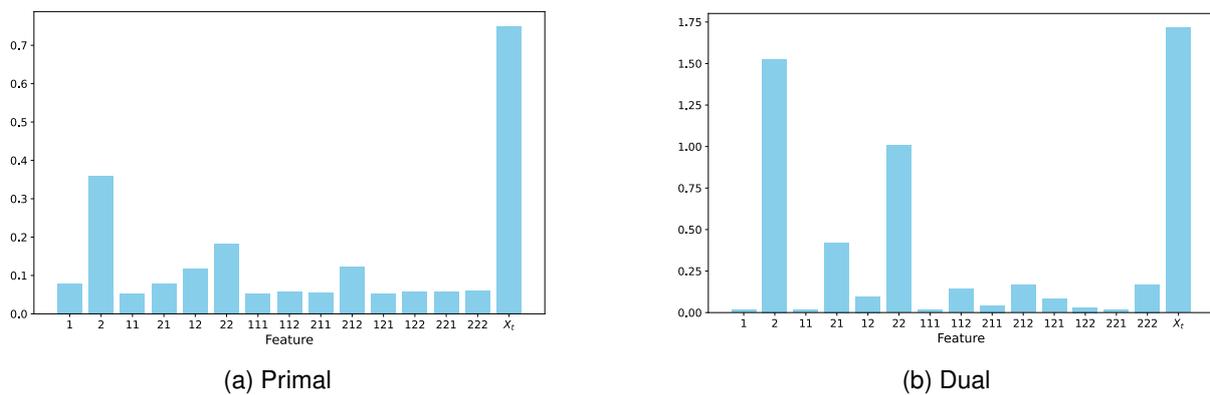


Figure 5: Feature importance in the rough Bergomi model with  $H = 0.07$  and  $\rho = -0.9$ .

tance of different entries of the level-3 signature in the deep-signature pricing methodologies. This is simply done by first training the primal and dual model, and then for each fixed feature, we shuffle the corresponding samples and measure the error between the point-estimate and the shuffled-estimate.

This procedure we repeat independently 20 times for each feature and in Figures 5 we plot the averaged error errors. Notice that the higher this value is, the more influence the corresponding feature has for computing the options price. In both methods, we observe that the "Markovian state features  $X_t$ ,  $v_t = \langle 2, \hat{V}_{0,t} \rangle$ ,  $v_t^2 = \langle 22, \hat{V}_{0,t} \rangle$  appear to have the biggest influence. Apart from that, we can observe that the integrals  $\langle 21, \hat{V}_{0,t} \rangle = \int_0^t v_u d\langle X \rangle_u$  (resp.  $\int_0^t v_u du$ ), and  $\langle 212, \hat{V}_{0,t} \rangle = \int_0^t \int_0^u v_r d\langle X \rangle_r dX_u$  (resp.  $\langle 212, \hat{V}_{0,t} \rangle = \int_0^t \int_0^u v_r dr dX_u$ ) seem to carry the most important information about the past of the volatility.

## A Proof deep-signature optimal stopping

**Proof of Proposition 3.3** The first step is to generalize the global approximation result [6, Theorem 2.8]. Leaving details to [6, Section 2], we are given a measure space  $(\mathcal{X}, \mathcal{B}(\mathcal{X}), \mu)$ , where  $\mathcal{X}$  is the space of stopped, time-augmented, geometric  $\alpha$ -Hölder rough paths ([6, Definition 2.1]),  $\mathcal{B}(\mathcal{X})$  the Borel  $\sigma$ -algebra, and  $\mu$  a measure fulfilling certain assumptions ([6, Assumption 2.6]). Thanks to [6, Theorem 2.8], for any  $f \in L^p(\mathcal{X}, \mu)$  and  $\epsilon > 0$ , we can find a linear functional on the signature  $f^\epsilon(\mathbf{x}|_{[0,t]}) = \langle l^\epsilon, \mathbb{X}_{0,t}^{\leq \infty} \rangle$ , such that  $\|f - f_n\|_{L^p} \leq \epsilon/2$ . But due to  $L^p$ -universality of DNNs, see [24, Proposition 2], denoting by  $k_\epsilon$  the truncation level of  $l^\epsilon$  in  $\mathcal{W}$ , for  $m$  large enough we can find  $(\beta, \tilde{\ell}) \in \mathbb{R}^m \times (\mathcal{W}^{\leq k_\epsilon})^m$  such that  $\|\theta^{(\beta, \tilde{\ell})}(\mathbf{x}|_{[0,t]}) - f^\epsilon\|_{L^p} \leq \epsilon/2$ . An application of the triangle inequality then proves that for any  $f \in L^p(\mathcal{X}, \mu)$ , we can find a sequence  $(\theta_n)_{n \in \mathbb{N}} \subseteq \text{DNN}_{\text{sig}}^\sigma$  such that  $\|f - \theta_n\|_{L^p} \rightarrow 0$  as  $n \rightarrow \infty$ .

Now in the primal case, it is in fact possible to show that for all  $0 \leq n \leq N$ , we have  $\mathbb{E}[Z_{\tau_n^{(m,k)}} | \mathcal{F}_{t_n}] \rightarrow \mathbb{E}[Z_{\tau_n} | \mathcal{F}_{t_n}]$  as  $(m, k) \rightarrow \infty$  in  $L^2$ , from which the result follows immediately. Using backward induction, one can notice that the result clearly holds true for  $n = N$ . Assuming that the claim holds for  $n + 1$ , we denote by  $\theta_n^{(m,k)}$  the element  $\theta^{(\beta_P^{*,n}, \ell_P^{*,n})}$  with respect to the minimizer  $(\beta_P^{*,n}, \ell_P^{*,n})$  in (8). We can estimate  $\|\mathbb{E}[Z_{\tau_n^{(m,k)}} - Z_{\tau_n} | \mathcal{F}_{t_n}]\|_{L^2}$  exactly as it was done [6, Appendix A.1] until equation (A.1), to reduce the claim to the convergence of

$$L^{(m,k)} = \|\theta^{(m,k)}(\mathbf{X}|_{[0,t_n]}) - \mathbb{E}[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}]\|_{L^2}.$$

There is, however, a subtle difference here compared to the conclusion used in the linear case: The space  $\text{DNN}_{\text{sig}}^\sigma$  is not convex, which means that  $\theta^{(m,k)}(\mathbf{X}|_{[0,t_n]})$  with respect to (8) is not necessarily the orthogonal projection of  $Z_{\tau_{n+1}^{(k,m)}}$ . We can still conclude by introducing the minimizer  $(\hat{\beta}_P^{*,n}, \hat{\ell}_P^{*,n})$  similar to (8), but replacing  $Z_{\tau_{n+1}^{(k,m)}}$  by  $\mathbb{E}[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}]$ . Then we have

$$\begin{aligned} L^{(m,k)} &\leq \|\theta^{(m,k)}(\mathbf{X}|_{[0,t_n]}) - \mathbb{E}[Z_{\tau_{n+1}^{(m,k)}} | \mathcal{F}_{t_n}]\|_{L^2} + \|\mathbb{E}[Z_{\tau_{n+1}^{(m,k)}} - Z_{\tau_{n+1}} | \mathcal{F}_{t_n}]\|_{L^2} \\ &\leq \|\hat{\theta}^{(m,k)}(\mathbf{X}|_{[0,t_n]}) - \mathbb{E}[Z_{\tau_{n+1}^{(m,k)}} | \mathcal{F}_{t_n}]\|_{L^2} + \|\mathbb{E}[Z_{\tau_{n+1}^{(m,k)}} - Z_{\tau_{n+1}} | \mathcal{F}_{t_n}]\|_{L^2} \\ &\leq \|\hat{\theta}^{(m,k)}(\mathbf{X}|_{[0,t_n]}) - \mathbb{E}[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}]\|_{L^2} + 2\|\mathbb{E}[Z_{\tau_{n+1}^{(m,k)}} - Z_{\tau_{n+1}} | \mathcal{F}_{t_n}]\|_{L^2} \\ &\leq \|\hat{\theta}^{(m,k)}(\mathbf{X}|_{[0,t_n]}) - \mathbb{E}[Z_{\tau_{n+1}} | \mathcal{F}_{t_n}]\|_{L^2} + 2\|\mathbb{E}[Z_{\tau_{n+1}^{(m,k)}} - Z_{\tau_{n+1}} | \mathcal{F}_{t_{n+1}}]\|_{L^2}, \end{aligned}$$

where we used the triangle inequality in the first and third inequality, the fact that  $\theta^{(m,k)}$  minimizes the  $L^2$  distance to  $\mathbb{E}[Z_{\tau_{n+1}^{(m,k)}} | \mathcal{F}_{t_n}]$  in the second inequality, and the contraction property of conditional expectations in the last inequality. Now the last term in the inequality converges by induction hypothesis, while the first one converges thanks to the global approximation result as  $m, k \rightarrow \infty$ .

Finally, let us outline also the proof of the dual case, which closely follows the techniques used in [6, Appendix A.2]. First, an application of the global approximation allows us to equivalently write (4) as

$$Y_0^N = \inf_{\theta \in \text{DNN}_{\text{sig}}^\sigma} \mathbb{E} \left[ \max_{0 \leq n \leq N} (Z_{t_n} - M_{t_n}^\theta) \right], \quad M_{t_n}^\theta = \int_0^{t_n} \theta(\mathbf{X}|_{[0,s]}) dW_s \quad (18)$$

Indeed, leaving the detailed technique to the proof of [6, Theorem 3.7], this follows from the following observation: As already discussed in the beginning of Section 2.2, since  $\mathcal{F}^X = \mathcal{F}^W$ , by martingale representation it is enough to minimize (4) over  $L^2$ -progressive processes, which in turn can be approximated arbitrarily well by  $\theta \in \text{DNN}_{\text{sig}}^\sigma$  thanks to the global approximation result. Next we define  $Y_0^{N,(m,k)} = \mathbb{E}[\max_{0 \leq n \leq N} (Z_{t_n} - M_{t_n}^{(m,k)})]$ , where we recall  $M^{(m,k)}$  was defined with respect to the minimizer (9), and note that  $Y_0^{N,(m,k)} \geq Y_0^N$ . Then, by (18), for every  $\epsilon > 0$ , there exists an element  $\theta^\epsilon \in \text{DNN}_{\text{sig}}^\sigma$  such that  $Y_0^N + \epsilon \geq \mathbb{E}[\max_{0 \leq n \leq N} (Z_{t_n} - M_{t_n}^{\theta^\epsilon})]$ , so that

$$0 \leq Y_0^{N,(m,k)} - Y_0^N \leq \epsilon + Y_0^{N,(m,k)} - \mathbb{E} \left[ \max_{0 \leq n \leq N} (Z_{t_n} - M_{t_n}^{\theta^\epsilon}) \right] \leq \epsilon,$$

where the last inequality holds for  $(m, k)$  large enough, such that  $\theta^\epsilon = \theta^{(\tilde{\beta}, \tilde{\ell})}$  for some  $(\tilde{\beta}, \tilde{\ell}) \in \mathbb{R}^m \times (\mathcal{W}^{\leq k})^m$ . □

**Remark A.1** *The same results remain true when working with the log-signature introduced in Remark 3.4. For the global approximation result, one can simply note that for any  $f \in L^p(\mathcal{X}, \mu)$ , we know that  $f \stackrel{L^p}{\approx} \langle \ell, \mathbb{X}^{\leq K} \rangle = \langle \ell, \exp^\otimes(\mathbb{L}^{\leq K}) \rangle$  for  $K$  large enough, where  $\exp^\otimes$  is the (continuous!) inverse of the  $\log^\otimes$  introduced in Remark 3.4. Using again [24, Proposition 2], we can approximate this exponential on the truncated log-signature by a DNN. In summary,  $f \stackrel{L^p}{\approx} \sum_j \beta_j \sigma(\langle \tilde{\ell}_j, \mathbb{L}^{< \infty} \rangle)$  in the sense of the global approximation result discussed in the beginning of the proof. Then, all the arguments can be repeated, relying on the set  $\text{DNN}_{\text{log}}^\sigma$ , which replaces the signature with the log-signature.*

## References

- [1] Peter Bank, Christian Bayer, Peter K Friz, and Luca Pelizzari. Rough PDEs for local stochastic volatility models. *arXiv preprint arXiv:2307.09216*, 2023.
- [2] Christian Bayer and Simon Breneis. Efficient option pricing in the rough Heston model using weak simulation schemes. *Quantitative Finance*, pages 1–15, 2024.
- [3] Christian Bayer, Peter Friz, and Jim Gatheral. Pricing under rough volatility. *Quantitative Finance*, 16(6):887–904, 2016.
- [4] Christian Bayer, Peter K Friz, Masaaki Fukasawa, Jim Gatheral, Antoine Jacquier, and Mathieu Rosenbaum. *Rough volatility*. SIAM, 2023.
- [5] Christian Bayer, Paul P Hager, Sebastian Riedel, and John Schoenmakers. Optimal stopping with signatures. *The Annals of Applied Probability*, 33(1):238–273, 2023.
- [6] Christian Bayer, Luca Pelizzari, and John Schoenmakers. Primal and dual optimal stopping with signatures. *arXiv preprint arXiv:2312.03444*, 2023.

- [7] Christian Bayer, Raúl Tempone, and Sören Wolfers. Pricing American options by exercise rate optimization. *Quantitative Finance*, 20(11):1749–1760, 2020.
- [8] Denis Belomestny. Solving optimal stopping problems via empirical dual optimization. *The Annals of Applied Probability*, 23(5):1988 – 2019, 2013.
- [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] Thomas Cass and Cristopher Salvi. Lecture notes on rough paths and applications to machine learning. *arXiv preprint arXiv:2404.06583*, 2024.
- [11] Kuo-Tsai Chen. Integration of paths, geometric invariants and a generalized Baker-Hausdorff formula. *Annals of Mathematics*, 65(1):163–178, 1957.
- [12] Ilya Chevyrev and Harald Oberhauser. Signature moments to characterize laws of stochastic processes. *The Journal of Machine Learning Research*, 23(1):7928–7969, 2022.
- [13] Vijay V Desai, Vivek F Farias, and Ciamac C Moallemi. Pathwise optimization for optimal stopping problems. *Management Science*, 58(12):2292–2308, 2012.
- [14] Petros Drineas, Michael W Mahoney, and Nello Cristianini. On the Nyström method for approximating a gram matrix for improved kernel-based learning. *journal of machine learning research*, 6(12), 2005.
- [15] Omar El Euch and Mathieu Rosenbaum. The characteristic function of rough Heston models. *Mathematical Finance*, 29(1):3–38, 2019.
- [16] Peter K Friz and Martin Hairer. *A course on rough paths*. Springer, 2020.
- [17] Peter K Friz and Nicolas B Victoir. *Multidimensional stochastic processes as rough paths: theory and applications*, volume 120. Cambridge University Press, 2010.
- [18] Ludovic Goudenège, 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.
- [19] Jasdeep Kalsi, Terry Lyons, and Imanol Perez Arribas. Optimal execution with rough path signatures. *SIAM Journal on Financial Mathematics*, 11(2):470–493, 2020.
- [20] Ioannis Karatzas and Steven Shreve. *Brownian motion and stochastic calculus*, volume 113. Springer Science & Business Media, 1991.
- [21] Patrick Kidger, Patric Bonnier, Imanol Perez Arribas, Cristopher Salvi, and Terry Lyons. Deep signature transforms. *Advances in Neural Information Processing Systems*, 32, 2019.
- [22] Bernard Lapeyre and Jérôme Lelong. Neural network regression for Bermudan option pricing. *Monte Carlo Methods and Applications*, 27(3):227–247, 2021.
- [23] Maud Lemerrier and Terry Lyons. A high order solver for signature kernels. *arXiv preprint arXiv:2404.02926*, 2024.

- [24] Moshe Leshno, Vladimir Ya Lin, Allan Pinkus, and Shimon Schocken. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural networks*, 6(6):861–867, 1993.
- [25] 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.
- [26] Terry J Lyons. Differential equations driven by rough signals. *Revista Matemática Iberoamericana*, 14(2):215–310, 1998.
- [27] Ryan McCrickerd and Mikko S Pakkanen. Turbocharging Monte Carlo pricing for the rough Bergomi model. *Quantitative Finance*, 18(11):1877–1886, 2018.
- [28] Goran Peskir and Albert Shiryaev. *Optimal stopping and free-boundary problems*. Springer, 2006.
- [29] Jeremy Reizenstein. Calculation of iterated-integral signatures and log signatures. *arXiv preprint arXiv:1712.02757*, 2017.
- [30] Jeremy Reizenstein and Benjamin Graham. The iisignature library: efficient calculation of iterated-integral signatures and log signatures. *arXiv preprint arXiv:1802.08252*, 2018.
- [31] Leonard CG Rogers. Monte Carlo valuation of American options. *Mathematical Finance*, 12(3):271–286, 2002.
- [32] Cristopher Salvi, Thomas Cass, James Foster, Terry Lyons, and Weixin Yang. The signature kernel is the solution of a Goursat PDE. *SIAM Journal on Mathematics of Data Science*, 3(3):873–899, 2021.
- [33] Bernhard Schölkopf, Ralf Herbrich, and Alex J Smola. A generalized representer theorem. In *International conference on computational learning theory*, pages 416–426. Springer, 2001.
- [34] Ingo Steinwart. *Support Vector Machines*. Springer, 2008.
- [35] Csaba Toth, Patric Bonnier, and Harald Oberhauser. Seq2tens: An efficient representation of sequences by low-rank tensor projections. In *International Conference on Learning Representations*, 2021.
- [36] Csaba Toth, Harald Oberhauser, and Zoltan Szabo. Random Fourier signature features. *arXiv preprint arXiv:2311.12214*, 2023.
- [37] Christopher Williams and Matthias Seeger. Using the Nyström method to speed up kernel machines. *Advances in neural information processing systems*, 13, 2000.