

A PATH-DEPENDENT PDE SOLVER BASED ON SIGNATURE KERNELS

ALEXANDRE PANNIER AND CRISTOPHER SALVI

ABSTRACT. We develop a provably convergent kernel-based solver for path-dependent PDEs (PPDEs). Our numerical scheme leverages signature kernels, a recently introduced class of kernels on path-space. Specifically, we solve an optimal recovery problem by approximating the solution of a PPDE with an element of minimal norm in the signature reproducing kernel Hilbert space (RKHS) constrained to satisfy the PPDE at a finite collection of collocation paths. In the linear case, we show that the optimisation has a unique closed-form solution expressed in terms of signature kernel evaluations at the collocation paths. We prove consistency of the proposed scheme, guaranteeing convergence to the PPDE solution as the number of collocation points increases. Finally, several numerical examples are presented, in particular in the context of option pricing under rough volatility. Our numerical scheme constitutes a valid alternative to the ubiquitous Monte Carlo methods.

1. INTRODUCTION

1.1. Path-dependent PDEs. The Feynman-Kac formula asserts that solutions to backward Kolmogorov equations have a probabilistic representation as $u_0(t, x) = \mathbb{E}[\phi(Y_T) | Y_t = x]$, for $0 \leq t \leq T$, $x \in \mathbb{R}^d$, where Y is an underlying Markov process. If Y models the evolution of a financial asset price, then this expectation corresponds to the price of a European-type derivative written on this asset. In its seminal paper, Dupire [21] introduces a first generation of path-dependent PDEs (PPDE) which solution is represented as $u_1(t, \gamma) = \mathbb{E}[\phi(Y_{[0,T]}) | Y_{[0,t]} = \gamma]$ where ϕ acts on the whole past of Y and γ is a path. This expectation represents the price of a path-dependent derivative. Rough volatility models, and more generally stochastic Volterra processes, are not Markov processes hence neither formulation applies. The recent framework of [52] tackles precisely this class of processes, for which the path-dependence is inherent to the dynamics. We illustrate this with an example to highlight the differences with the aforementioned cases. Let us consider a one factor rough volatility model

$$(1.1) \quad X_t = X_0 + \int_0^t \sqrt{\psi_s(\widehat{W}_s)} dB_s - \frac{1}{2} \int_0^t \psi_s(\widehat{W}_s) ds, \quad \widehat{W}_t := \int_0^t K(t, r) dW_r,$$

where B and W are correlated Brownian motions and generate their natural filtration $(\mathcal{F}_t)_{t \in [0, T]}$, the kernel K is square integrable and typically taken to be $K(t, r) = (t - r)^{H - \frac{1}{2}}$, $H \in (0, 1/2)$. In this situation, \widehat{W} is a Gaussian Volterra process, which is neither a semimartingale nor a Markov process, hence why classical tools are not suitable. For $s \geq t$, we make the orthogonal decomposition $\widehat{W}_s = \Theta_s^t + I_s^t$ where $\Theta_s^t = \int_0^t K(s, r) dW_r$ is \mathcal{F}_t -measurable while $I_s^t = \int_t^s K(s, r) dW_r$ is independent from \mathcal{F}_t .

UNIVERSITÉ PARIS CITÉ, LABORATOIRE DE PROBABILITÉS, STATISTIQUES ET MODÉLISATION
 IMPERIAL COLLEGE LONDON, DEPARTMENT OF MATHEMATICS AND IMPERIAL-X
E-mail addresses: pannier@lpsm.paris, c.salvi@imperial.ac.uk.

Date: March 19, 2024.

2020 *Mathematics Subject Classification.* 35R15, 60L10, 65N35, 91G60.

Key words and phrases. Path signature, kernel methods, path-dependent PDEs, rough volatility.

As a consequence one recovers the Markovian structure

$$\begin{aligned}
\mathbb{E}[\phi(X_T)|\mathcal{F}_t] &= \mathbb{E} \left[\phi \left(X_t + \int_t^T \sqrt{\psi_s(\widehat{W}_s)} dB_s - \frac{1}{2} \int_t^T \psi_s(\widehat{W}_s) ds \right) \middle| \mathcal{F}_t \right] \\
&= \mathbb{E} \left[\phi \left(X_t + \int_t^T \sqrt{\psi_s(\Theta_s^t + I_s^t)} dB_s - \frac{1}{2} \int_t^T \psi_s(\Theta_s^t + I_s^t) ds \right) \middle| \mathcal{F}_t \right] \\
&= \mathbb{E} \left[\phi \left(x + \int_t^T \sqrt{\psi_s(\gamma_s + I_s^t)} dB_s - \frac{1}{2} \int_t^T \psi_s(\gamma_s + I_s^t) ds \right) \middle| X_t = x, \Theta^t = \gamma_{[t,T]} \right] \\
&= \mathbb{E} [\phi(X_T)|X_t, \Theta_{[t,T]}^t],
\end{aligned}$$

The path-dependence of the option price is now clear and entails there is a measurable map u_2 such that $u_2(t, x, \gamma) = \mathbb{E} [\phi(X_T)|X_t = x, \Theta^t = \gamma_{[t,T]}]$. Building on Viens and Zhang's functional Itô formula, further works [10, 53] study the associated path-dependent PDE of second generation of the form

$$(1.2) \quad \begin{cases} \mathcal{L}u(t, x, \gamma) = 0, \\ u(T, x, \gamma) = \phi(x), \end{cases}$$

for all $t \in [0, T]$, $x \in \mathbb{R}$, $\gamma \in C([0, T]; \mathbb{R})$, and where the linear operator \mathcal{L} is a combination of derivatives of all inputs. The pathwise derivatives are defined in Section 2.2 while important well-posed examples can be found in Section 2.3.

Apart from the realm of rough volatility, stochastic Volterra processes have appeared in the modeling of electricity prices and turbulence [4, 38, 19, 3], in principal-agent problems [1], climate modeling [22, 50] and in the study of certain asymptotic regimes [27, 35, 26]. More broadly, fractional Brownian motions and related processes are present in a variety of situations where time-correlation is observed.

While the PPDE representation finds applications via studying its properties (e.g. weak error rates in [10]), this article is concerned with their numerical computation. The solution to a PPDE is a function of a path (and other finite dimensional inputs). Contrary to the wide literature on numerical schemes for high-dimensional PDEs, we are here interested in purely infinite dimensional inputs that we will treat as such instead of discretising them. Despite the central place of such PPDEs in mathematical finance, numerical methods pertaining to them are scarce, see Section 1.6 for more details. Our motivation is then twofold: design a numerical solver for such PDEs and overcome the technical challenges of learning a function on paths.

1.2. Signature kernels. Kernel methods are a well-established class of algorithms that are the key components of several machine learning models such as support vector machines (SVMs). The central idea of kernel methods is to lift the (possibly unstructured) input data to a (possibly infinite dimensional) Hilbert space by means of a nonlinear *feature map* defined in terms of a kernel. The advantage of doing so lies in the fact that a non-linear optimisation tasks in the original space becomes linear once lifted to the feature space. Thus, the solution of the original optimisation problem can often be expressed entirely in terms of evaluations of the kernel on training points by means of celebrated representer theorems. The kernel can often be evaluated efficiently with no reference to the feature map, a property known as *kernel trick*. The selection of an effective kernel will usually be a task-dependent problem, and this challenge is particularly difficult when the data is sequential. Signature kernels [33, 44] are a class of *universal kernels* on sequential data which have received attention in recent years thanks to their efficiency in handling path-dependent problems [34, 47, 17, 45, 16, 30, 29, 37].

1.3. Kernels vs neural networks for PDEs. In recent years, PDE solvers based on neural networks and kernel methods have been introduced as alternatives to classical finite difference and finite elements schemes. The most famous examples of neural PDE solvers are the Deep Galerkin model (DGM) [49] and Physics informed neural networks (PINNs) [41]. These techniques essentially parameterise the

solution of the PDE as a neural network which is then trained using the PDE and its boundary conditions as loss function. More recently, DGM and PINNs have been extended to path-dependent PDEs by [43, 48, 31]. PDE solvers based on kernel methods have been investigated in various works [18, 12, 13], but the development of kernel methods for path-dependent PDEs has remained open. PDE solvers based on kernel methods hold potential for considerable advantages over neural PDE solvers, both in terms of theoretical analysis and numerical implementation. In effect, the theoretical analysis of neural PDE solvers is often limited to density/universal approximation results aimed at showing the existence of a network of a requisite size achieving a certain error rate, without guaranteeing whether this network is computable in practice via stochastic gradient descent. On the contrary, as we will demonstrate in this work, kernel-based PDE solvers are provably convergent and amenable to rigorous numerical analysis based on a rich arsenal of functional analytic results from the kernel literature.

1.4. Monte Carlo methods for pricing under rough volatility. The notion of *rough volatility* has recently disrupted stochastic modelling in quantitative finance [7], leaving no one indifferent. In this setting, the instantaneous volatility of assets is modelled as a fractional Brownian motion with small Hurst parameter as in (1.1) or as a stochastic Volterra process. This feature has been shown empirically to be consistent with historical market trajectories and to capture important stylised facts of equity options in a parsimonious way. However, this new paradigm comes with an increased computational cost compared to classical techniques. With the notable exception of the rough Heston model, the absence of Markovianity of volatility prevents any pricing tools other than Monte Carlo simulations. However, the new promises—for estimation and calibration—of these rough volatility models have encouraged deep and fast innovations in numerical methods for pricing, in particular the now standard Hybrid scheme [8]. The complexity of a Monte Carlo scheme scales roughly as $\mathcal{O}(h^{-2}Nd^2)$, where N is the number of sample paths, h is the discretisation step and d is the dimension of the process. This complexity can be further reduced to $\mathcal{O}((h \log(h))^{-1}Nd^2)$ with hybrid schemes [8]. In this paper, we propose a simulation-free alternative to Monte Carlo pricing constructing a kernel-based solver for the corresponding PPDE. The complexity of the resulting solver scales as $\mathcal{O}(h^{-2}\tilde{N}d^2)$ where this time \tilde{N} is the number of collocation paths. This cost can be further reduced to $\mathcal{O}(h^{-1}\tilde{N}d^2)$ if the operations are carried out on a GPU, see [44] for details. As our numerical experiments demonstrate, a transparent comparison of complexities between our approach and Monte Carlo for pricing under rough volatility can only be achieved by obtaining error rates for the convergence of the approximations, which we intend to explore as future work.

1.5. Contributions. In this paper we design a novel numerical solver for PPDEs based on signature kernel and accompanied by convergence guarantees. Our numerical scheme does not use any classical probabilistic representation or any other pre-existing solver; it constitutes a valid alternative to the ubiquitous Monte Carlo methods, in particular in the context of option pricing under rough volatility. The hurdles in the theoretical analysis are many-fold and their resolution contributes to the development of the theory of signature kernels. Challenges include the translation of the PPDE framework into one amenable for the use of signature kernels, the proof of existence and uniqueness of a global minimum in the optimisation problem as well as an analytic form of the solution by means of a representer theorem (Theorem 4.2 and Corollary 4.4), and a consistency result based on the adaptation of convergence results from [12] to this infinite-dimensional setting via a compactness argument for functions on path spaces (Theorem 4.6 and Lemma 4.8). Finally, several numerical examples are presented in Section 5 and the code is open-sourced at <https://github.com/crispitaagorico/sigppde>.

1.6. Related literature. Several deep learning and signature-based approaches have been proposed in recent years. In [43], the authors use a probabilistic approach based on the martingale representation theorem to design an objective function for solving time-discretised PPDEs via a parametric model combining RNNs and signatures. The martingale representation theorem is also central in [23], where the authors make use of a continuous-time RNNs developed in [39]. In both these papers, the signature

has to be truncated for the implementation. The authors of [48] extend the deep Galerkin framework proposed by [49] using an LSTM architecture at each time step. The only paper applicable to Volterra processes [31] develops a backward discretisation algorithm using neural networks, drawing inspiration from the deep learning methodology introduced in [28]. We note that none of the aforementioned papers present theoretical results on the proposed numerical schemes and they all rely on the training of neural networks to approximate the solution. Aloof from the learning paradigm, convergence guarantees of monotone schemes can be found in [42] for the viscosity solutions to first generation PPDEs. The present paper is, to our knowledge, the first to propose a numerical scheme for second generation PPDEs with convergence guarantees.

1.7. Outline. Section 2 describes the rigorous framework of PPDEs and how we intent to approximate them. This is followed in Section 3 by an introduction to signature kernels and the definitions we will need in the remainder of the paper. The main theoretical analysis and results are presented in Section 4 while the numerical experiments can be found in Section 5. We give an outlook on future research in Section 6. Finally, the Appendix A gathers some technical proofs.

2. PATH-DEPENDENT PDES

As the example from the introduction suggests, in situations where the underlying process is fractional, such as fractional Brownian motion, rough volatility models or stochastic Volterra processes, conditional expectations are functions of paths. In this section, we will present several instances where these functions are solutions to well-posed path-dependent PDEs. Beforehand, we must set the mathematical scene in a rigorous way. We introduce the notations and definitions needed to formalise equations such as (1.2). This framework enables to present central examples in subsection 2.3, then in subsection 2.4 we adapt this setup to the signature realm and finally outline the numerical method is subsection 2.5.

2.1. Path spaces. Let us set $d, e \in \mathbb{N}_0 := \mathbb{N} \cup \{0\}$, a finite time horizon $T > 0$ and an interval $\mathbb{T} := [0, T]$. Throughout this paper we deal with continuous paths $\gamma : \mathbb{T} \rightarrow \mathbb{R}^e$ for which we need to introduce a number of spaces and norms. When the interval \mathbb{T} and state space \mathbb{R}^e are clear from the context we will often omit them. The space of continuous functions from \mathbb{T} to \mathbb{R}^e , denoted $\mathcal{C}(\mathbb{T}, \mathbb{R}^e)$, is equipped with the supremum norm $\|x\|_{0;\mathbb{T}} := \sup_{t \in \mathbb{T}} |\gamma_t|$, unless stated otherwise. The notation $|\cdot|$ refers to the Euclidean norm in one or multiple dimensions. The space of càdlàg paths from \mathbb{T} to \mathbb{R}^e is denoted $D(\mathbb{T}, \mathbb{R}^e)$ while for all $t \in \mathbb{T}$ we will make use of the space of càdlàg paths which are continuous on $[t, T]$, that is

$$\mathcal{C}_t := \left\{ \gamma \in D(\mathbb{T}, \mathbb{R}^e) : \gamma|_{[t, T]} \in \mathcal{C}([t, T], \mathbb{R}^e) \right\}.$$

This space allows to consider paths of the form $\gamma + \varepsilon \eta \mathbb{1}_{[t, T]}$ where $\gamma, \eta \in \mathcal{C}(\mathbb{T}, \mathbb{R}^e) = \mathcal{C}_0$ and $\varepsilon > 0$ which in turn permits the definition of pathwise derivatives on this space. In parallel, for $p \geq 1$, we also leverage the p -variation seminorm which is defined for $0 \leq t < t' \leq T$ and a path $\gamma : [t, t'] \rightarrow \mathbb{R}^e$ as

$$|\gamma|_{p-var;[t, t']} := \left(\sup_{(t_i)_{i \in \mathcal{D}}} \sum_i |\gamma_{t_{i+1}} - \gamma_{t_i}|^p \right)^{1/p},$$

where \mathcal{D} is the set of all partitions of the form $t = t_0 < t_1 < \dots < t_n = t'$ for some $n \in \mathbb{N}$. The associated norm is $\|\gamma\|_{p-var;[t, t']} = |\gamma_t| + |\gamma|_{p-var;[t, t']}$. This gives rise to the spaces $C^{p-var}([t, t'], \mathbb{R}^e)$ of continuous paths with one strictly monotone coordinate¹ and finite p -variation norm. The case $p = 1$

¹This is a technical assumption needed to ensure injectivity of the signature and consequently the point separating property assumption needed to establish universality of the signature kernel by Stone-Weierstrass arguments. More generally one could work with equivalence classes obtained by quotienting \mathcal{C} by the so-called *tree-like equivalence relation* $x \sim_\tau y \iff S(x) = S(y)$; with the assumption that the paths have one strictly monotone coordinate, each equivalence class collapses to a singleton, which is the path itself. The monotone coordinate is usually taken to be time.

is called the bounded variation space and also denoted $BV([t, t'], \mathbb{R}^e)$. Of particular interest to us are the spaces $C_t^{p-var} := C^{p-var}([t, T], \mathbb{R}^e)$ for some $t \in \mathbb{T}$. While the path space \mathcal{C}_t in the original theory of [52] allows for one jump, this would lead to severe complications for the application of signature kernels. We show in the next subsection how to adapt our framework to avoid the introduction of jumps. Furthermore, we want to equip the space $BV_0 := C_0^{1-var}$ with a topology that renders the signature map continuous (see Section 3 for the details). We could choose any $1 \leq p$ -variation topology for this task, however the necessity to characterise compact sets of BV_0 later on forces us to consider $p > 1$. This task involves the introduction of Hölder spaces $C^\alpha([t, t'], \mathbb{R}^e)$, with $\alpha \in (0, 1)$ endowed with the α -Hölder norm

$$\|\gamma\|_{\alpha;[t,t']} := |\gamma_t| + |\gamma|_{\alpha;[t,t']} := |\gamma_t| + \sup_{0 \leq s < t \leq T} \frac{|\gamma_t - \gamma_s|}{|t - s|^\alpha}.$$

If the interval $[t, t']$ is clear from the context, for instance if $\gamma \in BV([t, T], \mathbb{R}^e)$, we omit to write it as a subscript. Having in mind to characterise compact subsets of C^{p-var} , we follow [20] and intersect it with a Hölder space: let $C^{p-var, \alpha}([t, t'], \mathbb{R}^e) := C^{p-var}([t, t'], \mathbb{R}^e) \cap C^\alpha([t, t'], \mathbb{R}^e)$ be the space of α -Hölder continuous paths with finite p -variation. The corresponding norm is $\|\gamma\|_{p-var, \alpha;[t,t']} := |\gamma_t| + |\gamma|_{\alpha;[t,t']} + |\gamma|_{p-var;[t,t']}$. It is then proven in [20, Theorema A.8] that the embedding $C^{p-var, \alpha} \hookrightarrow C^{p'-var, \alpha'}$ is compact for all $1 \leq p < p' \leq \infty$ and $0 \leq \alpha' < \alpha < 1$ such that $\alpha p < 1$ and $\alpha' p' < 1$. In particular, bounded subsets of $C^{1-var, \alpha}$ are compact with respect to the p -variation topology for all $\alpha \in (0, 1)$ and $p > 1$, see also [25, Proposition 5.30].

The solutions to the path-dependent PDEs of interest are functionals of time, space and paths, hence we define suitable space and distance, borrowed from [52]:

$$\tilde{\Omega} = \left\{ (t, x, \gamma) \in \mathbb{T} \times \mathbb{R}^d \times \mathcal{C}_t \right\}, \quad \tilde{\mathbf{d}}((t, x, \gamma), (\bar{t}, \bar{x}, \bar{\gamma})) := |t - \bar{t}| + |x - \bar{x}| + \|\gamma - \bar{\gamma}\|_{0; \mathbb{T}}.$$

Note that γ lives in \mathcal{C}_t which depends directly on t . We will denote by Ω the state space with BV_0 instead of \mathcal{C}_t , that is

$$\Omega := \left\{ (t, x, \gamma) \in \mathbb{T} \times \mathbb{R}^d \times BV_0 \right\}.$$

The analogous distance is then

$$\mathbf{d}((t, x, \gamma), (\bar{t}, \bar{x}, \bar{\gamma})) := |t - \bar{t}| + |x - \bar{x}| + \|\gamma - \bar{\gamma}\|_{p-var; \mathbb{T}}.$$

Notice that if B is a compact subset of BV_0 with respect to the p -variation topology then for any reals $a < b$, the set $\mathbb{T} \times [a, b]^d \times B$ is a compact subset of Ω with respect to \mathbf{d} .

2.2. Functions on continuous paths. This subsection presents the framework of [52]. Let $\mathcal{C}(\tilde{\Omega}) := \mathcal{C}(\tilde{\Omega}, \mathbb{R})$ denote the set of real-valued functions $u : \tilde{\Omega} \rightarrow \mathbb{R}$ continuous under $\tilde{\mathbf{d}}$. For $u \in \mathcal{C}(\tilde{\Omega})$, define the right time derivative

$$\partial_t u(t, x, \gamma) := \lim_{\varepsilon \downarrow 0} \frac{u(t + \varepsilon, x, \gamma) - u(t, x, \gamma)}{\varepsilon},$$

for all $t \in [0, T)$, provided the limit exists. By convention we set $\partial_t u(T, x, \gamma) = 0$. The derivative $\partial_x u(t, x, \gamma)$ is defined in the same way for all $(t, x, \gamma) \in \tilde{\Omega}$. We also introduce the directional derivative $\tilde{\partial}_\gamma^\eta u(t, x, \gamma)$ with respect to $\gamma \in \mathcal{C}_0$, in the direction of $\eta \in \mathcal{C}([t, T], \mathbb{R}^e)$:

$$(2.1) \quad u(t, x, \gamma + \varepsilon \eta \mathbb{1}_{[t, T]}) - u(t, x, \gamma) = \tilde{\partial}_\gamma^\eta u(t, x, \gamma) + o(\|\eta \mathbb{1}_{[t, T]}\|_{\mathbb{T}}).$$

Notice that the path is only perturbed on the interval $[t, T]$, hence the necessity of a path space allowing for jumps. The operator $\eta \mapsto \tilde{\partial}_\gamma^\eta u$ is linear on $\mathcal{C}([t, T], \mathbb{R}^e)$ and a Fréchet derivative; moreover if it exists it is equal to the Gateaux derivative. We observe crucially that both the Fréchet and Gateaux derivative are operators acting on functions from $\tilde{\Omega}$ to \mathbb{R} ; in particular the time dependence is essential. We define higher derivatives as linear operators on $(\mathcal{C}([t, T], \mathbb{R}^e))^n$, in particular the second derivative is a bilinear operator on $\mathcal{C}([t, T], \mathbb{R}^e) \times \mathcal{C}([t, T], \mathbb{R}^e)$ and we denote it $\tilde{\partial}_\gamma^{\eta \bar{\eta}} u(t, x, \omega)$ for $\eta, \bar{\eta} \in \mathcal{C}([t, T], \mathbb{R}^e)$.

Let $n, n_1, n_2, n_3 \in \mathbb{N}_0$. We denote the higher derivatives $\partial_t^{n_1}$, $\partial_x^{n_2}$ and $\tilde{\partial}_\gamma^\eta$, for $\eta = (\eta_1, \dots, \eta_{n_3}) \in (\mathcal{C}([t, T], \mathbb{R}^e))^{n_3}$; the order of the directional derivative is implied by the dimension of η . We say that a function $u \in \mathcal{C}(\tilde{\Omega})$ belongs to $\mathcal{C}^\mathcal{N}(\tilde{\Omega})$ if the linear operator $\partial_t^{n_1} \partial_x^{n_2} \tilde{\partial}_\gamma^\eta$ exists and is continuous with respect to $\tilde{\mathbf{d}}$ for all $(n_1, n_2, n_3) \in \mathcal{N} \subset \mathbb{N}_0^3$. The examples of interest arise from Kolmogorov equations, hence only the first order time derivative appears and is never combined with the spatial and pathwise derivatives, while the latter are present up to order two. The corresponding subset is $\mathcal{N} = \{(n_1, n_2, n_3) \in \mathbb{N}_0^3 : n_1 \leq 1, n_2 = n_3 = 0, \text{ or } n_1 = 0, n_2 + n_3 \leq 2\}$.

We intend to deal with fractional processes with singular kernels of the form $K(s, t) = (s-t)^{H-\frac{1}{2}} \mathbb{1}_{s>t}$ with $H \in (0, \frac{1}{2})$. Taken as a function of $s \in [t, T]$, these kernels turn out to be the direction of the pathwise derivative but lie outside of $\mathcal{C}(\mathbb{T}, \mathbb{R})$, hence an approximation procedure is required. For $0 \leq t < s \leq T$, we define the path $K^t(s) := K(s, t)$ and its truncated version $K^{\delta, t}(s) := K(s \wedge (t + \delta), t)$, for $\delta > 0$. The directional derivative for such a singular path is then defined, if it exists, as

$$(2.2) \quad \tilde{\partial}_\gamma^{K^t} u(t, x, \gamma) := \lim_{\delta \downarrow 0} \tilde{\partial}_\gamma^{K^{\delta, t}} u(t, x, \gamma),$$

and similarly for the higher derivatives. The generalisation of the space $\mathcal{C}^\mathcal{N}$ to this type of derivative is more involved and requires regularity estimates tailored to the speed of explosion of the kernel. For precise definitions, we will point towards the appropriate references in the examples below. We are now ready to state the equations we wish to solve.

2.3. The target PDEs. This paper's focus is on approximating functionals u which only depend on the path over $[t, T]$, that is there exist \tilde{u} such that $u(t, x, \gamma) = \tilde{u}(t, x, \gamma|_{[t, T]})$. This is the case of all the examples presented in this section. These functionals are solutions to linear PPDEs taking the general form

$$(2.3) \quad \tilde{\mathcal{L}}_t u(t, x, \gamma) = f(t, x, \gamma), \quad u(T, x, \gamma) = \phi(T, x, \gamma),$$

where $f, \phi : \tilde{\Omega} \rightarrow \mathbb{R}$ are measurable functions and the linear operator $\tilde{\mathcal{L}}_t$ is a combination of derivatives of all arguments. More precisely, there exist $\mathcal{N} \subset \mathbb{N}_0^3$, $a_n \in \mathbb{R}$ for all $n \in \mathcal{N}$ and $\eta \in (\mathcal{C}([t, T], \mathbb{R}^e))^{n_3}$ such that

$$(2.4) \quad \tilde{\mathcal{L}}_t := \sum_{n=(n_1, n_2, n_3) \in \mathcal{N}} a_n \partial_t^{n_1} \partial_x^{n_2} \tilde{\partial}_\gamma^\eta,$$

where the sum is taken over multi-indices $n \in \mathcal{N}$. We emphasise on the t dependence which is important because the directional derivative is a linear operator on $\mathcal{C}([t, T], \mathbb{R}^e)$. We need this notation instead of the more standard multi-index one in order to emphasise on the order of each derivative (in our applications $n_1 \leq 1$, $n_2 + n_3 \leq 2$) and also on the direction η .

More precisely, this paper will focus on three examples of parabolic second order PDEs which arise from the Feynman-Kac formula. The solution to these PPDEs thus have a probabilistic representation, a crucial advantage to assess the accuracy of our method. However we emphasise that the numerical scheme does not use this a priori knowledge as in supervised learning tasks. Much like deep Galerkin or PINN algorithms [49, 41], our method can be deployed to solve PPDEs for which no numerical method is known.

Our examples take as underlying a Gaussian Volterra process \widehat{W} defined as $\widehat{W}_t := \int_0^t K(t, r) dW_r$ for some kernel $K \in L^2(\mathbb{T}^2, \mathbb{R})$. The key object, already introduced in the introduction, is the double index process, for $0 \leq t \leq s$, $\Theta_s^t = \mathbb{E}_s[\widehat{W}_t] = \int_0^t K(s, r) dW_r$. In alignment with the prolific rough volatility literature, we will mostly focus on the specification $K(t, r) = (t-r)^{H-1/2}$ for $H \in (0, 1)$ but other choices are possible, see for instance [10, Assumption 2.6]. The current setup (introduced in [52]), as well as our numerical method, allow for path-dependent payoffs. Since \widehat{W} is a fractional process, the path-dependence is inherent and transfers to the PDE even for state-dependent terminal conditions (payoffs). For clarity of exposition and because well-posedness results are scarcer in the

path-dependent case, we will however not consider them and stick to state-dependent payoffs in this paper.

In the three subsequent subsections, we present (i) the model (ii) the value function u and (iii) the PPDE, in that order. We point to the suitable references for well-posedness and for more details. We emphasise that all three examples are particular cases of [10, Proposition 2.14]; this result ensures, under general assumptions, existence and uniqueness of a classical solution to the PPDE associated to a Volterra process.

2.3.1. Fractional Brownian motion. In this first toy example we consider test functions of $(\widehat{W}_s)_{s \in [t, T]}$ which corresponds to the state space $\widetilde{\Omega}$ with $d = 0$ and $e = 1$. Let us define the function

$$(2.5) \quad u(t, \gamma) := \mathbb{E} \left[\phi(\widehat{W}_T) + \int_t^T f_s(\widehat{W}_s) ds \mid \Theta^t = \gamma \right]$$

$$(2.6) \quad = \mathbb{E} \left[\phi \left(\gamma_T + \int_t^T K(T, s) dW_s \right) + \int_t^T f_s \left(\gamma_s + \int_t^s K(t, r) dW_r \right) ds \right].$$

In particular, we observe that $u(t, \Theta^t) = \mathbb{E} \left[\phi(\widehat{W}_T) + \int_t^T f_s(\widehat{W}_s) ds \mid \mathcal{F}_t \right]$. Note that the path-dependence of the conditional expectation is contained in the path Θ^t because the payoff only acts on \widehat{W}_s for $s \in [t, T]$, whereas a payoff depending on \widehat{W}_s for all $s \in [0, T]$ would require to consider the concatenated path $\widehat{W}_s \mathbb{1}_{s < t} + \Theta_s^t \mathbb{1}_{s \geq t}$.

Based on [52, Theorem 4.1], under some standard regularity assumptions on ϕ, ψ and K , u is a solution to the path-dependent heat equation

$$(2.7) \quad \partial_t u(t, \gamma) + \frac{1}{2} \widetilde{\partial}_\gamma^{K^t K^t} u(t, \gamma) + \psi(t, \gamma) = 0, \quad u(T, \gamma) = \phi(\gamma_T), \quad \text{for any } (t, \gamma) \in \widetilde{\Omega}.$$

It can also be proved [10] that this solution is unique in the space where the functional Itô formula [52, Theorem] is applicable.

2.3.2. VIX options under rough Bergomi model. Options on VIX and realised variance form an important class of derivatives, which can be seen as options with path-dependent payoffs. We define $\text{VIX}_T^2 = \frac{1}{\Delta} \int_T^{T+\Delta} \mathbb{E}_T[\psi_s(\widehat{W}_s)] ds$ where Δ is a positive constant corresponding to 30 days, and \widehat{W} is a Gaussian Volterra process in \mathbb{R}^e for any $e \in \mathbb{N}$. It is shown in [40], that an option on the VIX (future) has the representation

$$\mathbb{E} \left[\phi(\text{VIX}_T^2) \mid \mathcal{F}_t \right] = \mathbb{E} \left[\phi \circ \mathfrak{F} \left(\left\{ \Theta_s^t + J_s^{t, T} : s \in [T, T + \Delta] \right\} \right) \mid \Theta^t \right],$$

where $J_s^{t, T} = \Theta_s^T - \Theta_s^t$ is independent of \mathcal{F}_t and $\mathfrak{F}(\gamma) := \frac{1}{\Delta} \int_T^{T+\Delta} \mathbb{E}[\psi_s(\gamma_s + I_s^T)] ds$ where $I_s^T \sim \mathcal{N}(\mathbf{0}_d, \int_0^s K(T, r) K(T, r)^\top dr)$. Under natural assumptions on ϕ, ψ and K , the map $u(t, \gamma) := \mathbb{E} \left[\phi \circ \mathfrak{F} \left(\left\{ \gamma_s + J_s^{t, T} : s \in [T, T + \Delta] \right\} \right) \right]$ is the unique solution to the PPDE

$$(2.8) \quad \begin{cases} \partial_t u(t, \gamma) + \frac{1}{2} \widetilde{\partial}_\gamma^{K^t K^t} u(t, \gamma) = 0, & (t, \gamma) \in \Lambda, \\ u(T, \gamma) = \phi \circ \mathfrak{F}(\gamma). \end{cases}$$

2.3.3. Options under rough Bergomi model. In this last example, which was already presented in the introduction, the underlying is the (log-)asset price in a rough Bergomi-type model. The value function in the rough Bergomi model is indexed on the state space $\widetilde{\Omega}$ with $d = 1$. As we have already seen,

$$u(t, x, \gamma) := \mathbb{E} \left[\phi(X_T) \mid X_t = x, \Theta^t = \gamma \right],$$

which entail $u(t, X_t, \Theta^t) = \mathbb{E}[\phi(X_T) | \mathcal{F}_t]$. Based on [10, Theorem 2.25], natural assumptions on ϕ, ψ and K yield the well-posedness of the PPDE

$$(2.9) \quad \partial_t u + \zeta \psi_t(\gamma_t) \partial_x u + \frac{1}{2} \psi_t(\gamma_t) \partial_x^2 u + \frac{1}{2} \tilde{\partial}_\gamma^{K^t K^t} u + \rho \sqrt{\psi_t(\gamma_t)} \tilde{\partial}_\gamma^{K^t} (\partial_x u) = 0,$$

for all $(t, x, \gamma) \in \tilde{\Omega}$ and with boundary condition $u(T, x, \gamma) = \phi(x)$. Notice the boundary condition only applies to x and is not path-dependent.

Remark 2.1. In the rough Bergomi model, the paths Θ have a financial interpretation as they are linked to the forward variance curves ξ , a quantity (more or less) observed on the markets. We refer to [40, Remark 4.7 (5)] for more details.

2.3.4. An important remark on the singularity of the kernel. We emphasise that the PPDEs presented in these examples fit in (2.3) albeit with a non-continuous direction K^t . The pathwise derivative is thus defined as the limit of the regularised derivative with the approximation $K^{t, \delta}$ as in (2.2). One can show thanks to the probabilistic representations that the solution to the regularised PPDE converges to the solution of the singular one as δ goes to zero. Our numerical scheme approximates the PPDE (2.3) with a regularised kernel, for a fixed $\delta > 0$.

2.4. Functions on bounded variation paths. Making the framework of subsection 2.2 amenable to the application of signature kernels requires some carefulness. Firstly, the latter are defined for bounded variation paths, a subset of the continuous paths considered in the previous subsection. In this subsection, functions are thus defined over Ω instead of $\tilde{\Omega}$.

Further, we make the observation that all three examples presented above are functionals $(t, x, \gamma) \mapsto u(t, x, \gamma)$ that, for all $t \in \mathbb{T}$, only act on $(\gamma_s)_{s \in [t, T]}$, the path after time t . We will therefore restrict our domain of interest and rely crucially on this property. We denote by $\mathcal{C}(\Omega)$ the set of such continuous functions $u : \Omega \rightarrow \mathbb{R}$; they are characterised by the existence of a family $(u_{t,x})_{t \in \mathbb{T}, x \in \mathbb{R}^d}$ of maps from BV_t to \mathbb{R} such that $u(t, x, \gamma) = u_{t,x}(\gamma|_{[t, T]})$, and such that u is continuous with respect to \mathbf{d} . Note that this distance looks into the path on $[0, t]$, indeed without this information one would lose the positivity for two paths distinct on $[0, t]$, while considering only continuity of $u_{t,x}$ leads to issues in defining compact sets of Ω . As the signature is not continuous with respect to the \mathcal{C}_0 norm, the Fréchet derivative (2.1) is inappropriate. We thus have to choose a weaker form of pathwise derivative: the Gateaux derivative is consistent with the derivative of the signature kernel introduced in [45]. This is where the restriction to $[t, T]$ becomes useful and allows us to avoid jumps. Recall that the Gateaux derivative perturbs the path on the whole interval where it is defined

$$(2.10) \quad \partial_\gamma^\eta u(t, x, \gamma) = \lim_{\varepsilon \downarrow 0} \frac{u(t, x, \gamma + \varepsilon \eta) - u(t, x, \gamma)}{\varepsilon};$$

we instead define the pathwise derivative of $u \in \mathcal{C}(\Omega)$ in the direction $\eta \in \text{BV}_0$ as

$$(2.11) \quad \hat{\partial}_\gamma^\eta u(t, x, \gamma) := \partial_\gamma^{\eta|_{[t, T]}} u_{t,x}(\gamma|_{[t, T]}) := \lim_{\varepsilon \rightarrow 0} \frac{u_{t,x}((\gamma + \eta)|_{[t, T]}) - u_{t,x}(\gamma|_{[t, T]})}{\varepsilon}.$$

Note that this amounts to three different notions of pathwise derivatives $\tilde{\partial}, \hat{\partial}, \partial$. Similarly to the former setup in subsection 2.2, for $n, n_1, n_2, n_3 \in \mathbb{N}_0$, we can define the higher derivatives $\partial_t^{n_1}, \partial_x^{n_2}$ and $\hat{\partial}_\gamma^{n_3}$, for $\boldsymbol{\eta} = (\eta_1, \dots, \eta_{n_3}) \in (\text{BV}_0)^{n_3}$. We denote $\mathcal{C}^\mathcal{N}(\Omega)$ the space of functions admitting continuous such pathwise derivatives of order $\mathcal{N} \subset \mathbb{N}_0^3$. We can now define PPDEs on Ω similar to (2.3)

$$(2.12) \quad \mathcal{L}_t u(t, x, \gamma) = f(t, x, \gamma), \quad u(T, x, \gamma) = \phi(T, x, \gamma),$$

where $f, \phi : \Omega \rightarrow \mathbb{R}$ are measurable functions and the linear operator \mathcal{L}_t is a combination of derivatives of all arguments, analogue of $\tilde{\mathcal{L}}_t$ for functions independent of $\gamma_{[0, t]}$ albeit with a pathwise derivative of Gateaux type instead of Fréchet. More precisely, there exist $\mathcal{N} \subset \mathbb{N}_0^3$, $a_n \in \mathbb{R}$ for all $n \in \mathcal{N}$

and $\boldsymbol{\eta} \in (\text{BV}_0)^{n_3}$ such that

$$\mathcal{L}_t := \sum_{n=(n_1, n_2, n_3) \in \mathcal{N}} a_n \partial_t^{n_1} \partial_x^{n_2} \widehat{\partial}_{\boldsymbol{\eta}}^{n_3},$$

where the sum is taken over multi-indices $n \in \mathcal{N}$.

2.5. A kernel method for PPDEs. We saw that we cannot solve a PPDE of the type $\tilde{\mathcal{L}}_t u = f$ on $\tilde{\Omega}$ directly with signature kernel methods, hence why we defined an analogous weaker PPDE $\mathcal{L}_t u = f$ on $\Omega \subset \tilde{\Omega}$ which solution can be suitably represented with signature kernels. Functions admitting Fréchet derivatives also admit Gateaux derivatives and they are equal, therefore solutions to the former PPDE are also solutions to the latter. This will allow us in Theorem 4.6 to bridge the gap between the equation we solve and the equation we target. Now that we have this picture in mind, we can present the gist of our method, inspired by [12] and extended to the path-dependent case.

We consider a kernel κ , indexed on a compact subset \mathfrak{K} of Ω , as product of an RBF kernel and a signature kernel indexed on \mathbb{R}^{d+1} and BV_0 respectively. We call \mathcal{H} its associated RKHS. We pick $m+n$ collocation points $\omega^1, \dots, \omega^{m+n} \in \mathfrak{K}$ of the form $\omega^i = (t^i, x^i, \gamma^i)$ and order them in such a way that $0 \leq t^i < T$ for $i = 1, \dots, m$ and $t^i = T$ for $i = m+1, \dots, m+n$. The former belong to the interior of the domain while the latter correspond to the boundary (terminal) condition. The main idea consists in approximating the solution of the PPDE (2.12) with a minimiser of the following optimal recovery problem

$$(2.13) \quad \min_{u \in \mathcal{H}} \frac{1}{2} \|u\|_{\mathcal{H}}^2 \quad \text{such that} \quad \begin{cases} \mathcal{L}_t u(\omega^i) = f(\omega^i) & i = 1, \dots, m \\ u(\omega^i) = \phi(\omega^i) & i = m+1, \dots, m+n. \end{cases}$$

In other words, we search for the function with minimal RKHS norm that satisfies the PDE constraints at all the collocation points. This problem formulation raises a number of questions, of both theoretical and practical importance:

- (1) Does the domain of \mathcal{L}_t contain \mathcal{H} ?
- (2) Does the minimisation problem (2.13) have a unique solution $u_{m,n}$?
- (3) How does one solve this optimisation over an infinite-dimensional space ?
- (4) Does the PPDE (2.3) have a unique classical solution u^* ?
- (5) Does $u_{m,n}$ converge to u^* as $m, n \rightarrow \infty$?

After designing our kernel and RKHS in Section 3, we answer positively to all these questions under appropriate conditions. In order to successively execute limit arguments, we exploit that \mathfrak{K} is a compact set of Ω , which we define explicitly under the topology induced by \mathbf{d} . Exploiting the robust signatures introduced in [14] may allow to lift this assumption; we will investigate this idea in the future. The interested reader may find a summary of the answers as follows.

- (1) Yes, Proposition 4.1 proves this assertion.
- (2) Yes, as this is a type of optimal recovery problem, see Theorem 4.2.
- (3) As detailed in Corollary 4.4, the optimal recovery problem reduces to a finite dimensional quadratic optimisation problem with linear constraints. The latter can be solved via gradient descent, aloof from sophisticated learning algorithms.
- (4) As pointed out in Section 2.3, existence and uniqueness of this PPDE holds under some conditions, which are thoroughly checked in the given references for the examples of interest.
- (5) Theorem 4.6 ensures convergence towards the true solution under two different sets of assumptions, both requiring that u^* is the unique solution of the PPDE (2.3) and that the collocation points form a countable dense set of \mathcal{H} .

Let us summarise the several levels of approximation. We numerically solve the optimisation problem (2.13), which is a discretisation of the PPDE (2.12) over Ω . Under the right set of assumptions, the solution to the optimisation problem actually approximates the solution to the PPDE (2.3) over $\tilde{\Omega}$. Finally, the latter is an approximation of the PPDE with a singular direction $\boldsymbol{\eta}$. While our setup is

designed to match PPDEs of second generation (arising from Volterra processes), it could be adapted to the first generation of Dupire (corresponding to path-dependent payoffs).

3. SIGNATURE KERNELS

Our numerical solver is based on so-called signature kernels, a special class of kernels indexed on BV. In this section we summarise their main properties of such signature kernels and highlight important aspects related to their numerical evaluation. We begin by recalling definition and properties of classical kernels.

3.1. Classical kernels. A scalar-valued kernel on a set Ω is a symmetric function of the form $\kappa : \Omega \times \Omega \rightarrow \mathbb{R}$. Such kernel is said to be positive semidefinite if for any $n \in \mathbb{N}$ and points $\omega_1, \dots, \omega_n \in \Omega$, the Gram matrix $\mathcal{K} := (\kappa(\omega_i, \omega_j))_{i,j}$ is positive semidefinite, i.e. if for any $\alpha_1, \dots, \alpha_n \in \mathbb{R}$

$$\sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j \kappa(\omega_i, \omega_j) \geq 0.$$

Recall that a Hilbert space \mathcal{H} of functions defined on a set Ω is a *reproducing kernel Hilbert space* (RKHS) over Ω if, for each $\omega \in \Omega$, the point evaluation functional at ω , $f \mapsto f(\omega)$, is a continuous linear functional, i.e. there exists a constant $C_\omega \geq 0$ such that

$$|f(\omega)| \leq C_\omega \|f\|_{\mathcal{H}}, \quad \text{for all } f \in \mathcal{H}.$$

The classical Moore-Aronszajn Theorem [2] ensures that if κ is a positive semidefinite kernel, then there exists a unique RKHS \mathcal{H} such that κ has the following *reproducing property*

$$(3.1) \quad \langle \kappa(\omega, \cdot), f \rangle_{\mathcal{H}} = f(\omega), \quad \text{for all } f \in \mathcal{H} \text{ and } \omega \in \Omega.$$

When dealing with questions related to function approximation, one is often interested in understanding how well elements of an RKHS can approximate a function in a given class. One important such class is the family of continuous functions, leading to the notion of *cc-universality* [51]. If the set Ω is equipped with an arbitrary topology and $\kappa : \Omega \times \Omega \rightarrow \mathbb{R}$ is a continuous, symmetric, positive semidefinite kernel on Ω with RKHS \mathcal{H} , then we say that κ is cc-universal on Ω if, for every compact subset $\mathcal{K} \subset \Omega$, \mathcal{H} is dense in $\mathcal{C}(\mathcal{K})$ in the topology of uniform convergence.

Kernels on Euclidean spaces such as linear, polynomial, Gaussian and Matérn kernels have been extensively studied in prior literature. Arguably, the most popular choice (which is also the one we make in our experiments) of cc-universal kernel on $\mathbb{R}_+ \times \mathbb{R}^d$ is the *radial basis function* (RBF) kernels $g_\sigma : \mathbb{R}_+ \times \mathbb{R}^d \times \mathbb{R}_+ \times \mathbb{R}^d \rightarrow \mathbb{R}$ defined for $(s, x), (t, y) \in \mathbb{R}_+ \times \mathbb{R}^d$ as

$$(3.2) \quad g_\sigma((s, x), (t, y)) = \exp\left(-\frac{\|(s, x) - (t, y)\|^2}{2\sigma^2}\right), \quad \sigma \in \mathbb{R}.$$

To handle solutions of PPDEs we need to introduce kernels indexed on BV. *Signature kernels* [33, 44] provide a natural class of cc-universal of kernels on BV as we shall discuss next.

3.2. Signature kernels. The definition of signature kernel require an initial algebraic setup. Denote by \otimes the standard tensor product of vector spaces. Let V be a Hilbert space with inner product $\langle \cdot, \cdot \rangle_V$. In this section, we adopt a more general notation, although we will specifically consider V as the vector space \mathbb{R}^e in the sequel of the paper. For any $n \in \mathbb{N}$, the canonical Hilbert-Schmidt inner product $\langle \cdot, \cdot \rangle_{V^{\otimes n}}$ on $V^{\otimes n}$ is defined as

$$(3.3) \quad \langle a, b \rangle_{V^{\otimes n}} = \prod_{i=1}^n \langle a_i, b_i \rangle_V$$

where $a = (a_1, \dots, a_n)$ and $b = (b_1, \dots, b_n)$ are tensors in $V^{\otimes n}$. Let $T((V)) = \prod_{n=1}^{\infty} V^{\otimes n}$ the vector space of formal tensor series and $T(V) \subset T((V))$ the space of formal tensor polynomials defined as

$$T(V) = \{a \in T((V)) \mid a \text{ has a finite number of non-zero entries}\}.$$

There are many ways of extending by linearity the inner product in (3.3) to an inner product on $T(V)$. The simplest way to achieve this is by defining the inner product as follows

$$(3.4) \quad \langle a, b \rangle_{T((V))} = \sum_{n=0}^{\infty} \langle a_n, b_n \rangle_{V^{\otimes n}}$$

for any $a = (a_0, a_1, \dots), b = (b_0, b_1, \dots)$ in $T(V)$.

Remark 3.1. A more general class of inner products was introduced in [11] by means of a weight function $\phi : \mathbb{N}_0 \rightarrow \mathbb{R}_+$ in the following way

$$(3.5) \quad \langle a, b \rangle_{T((V))} = \sum_{n=0}^{\infty} \phi(n) \langle a_n, b_n \rangle_{V^{\otimes n}}.$$

We denote by $\overline{T(V)}$ be the Hilbert space obtained by completing $T(V)$ with respect to $\langle \cdot, \cdot \rangle_{T((V))}$.

Supposing that V is d -dimensional, for some $d > 1$, it will be convenient to introduce the set of *words* \mathbb{W} in the non-commuting letters $\{1, \dots, d\}$, i.e. ordered m -tuples $\omega = (i_1, i_2, \dots, i_m) \in \{1, \dots, d\}^m$ with $|\omega| := m \in \mathbb{N}$.

The principal ingredient needed to define the signature kernel is a classical transform in stochastic analysis known as the *signature*. Next we recall its definition.

Let $\text{BV}(\mathbb{T}, V)$ be defined as in Section 2.1 where the norm in the definition of p -variation is the norm in V . For any closed interval $[t, t'] \subset \mathbb{T}$ and any path $\gamma \in \text{BV}(\mathbb{T}, V)$, the signature of γ over $[t, t']$ is the following infinite collection

$$S(\gamma)_{t,t'} = \left(S(\gamma)_{t,t'}^{(\omega)} \right)_{\omega \in \mathbb{W}} \in \overline{T(V)} \cong l^2(\mathbb{W})$$

of Riemann-Stieltjes iterated integrals

$$(3.6) \quad S(\gamma)_{t,t'}^{(\omega)} = \int_{t < t_1 < \dots < t_m < t'} d\gamma_{t_1}^{(i_1)} \dots d\gamma_{t_m}^{(i_m)}, \quad \text{for any } \omega = (i_1, \dots, i_m) \in \mathbb{W}.$$

When it is clear from the context, we will suppress the dependence on the interval and instead denote the signature of γ by $S(\gamma)$. We refer the interested reader to [32, 24] for an account and examples on the use of signature methods in machine learning.

The signature of a BV path γ can be equivalently defined as the solution of a $\overline{T(V)}$ -valued differential equations controlled by γ , as stated in the following lemma.

Lemma 3.2. *Let $\gamma \in \text{BV}(\mathbb{T}, V)$. Then, the linear controlled differential equation*

$$(3.7) \quad dY_s = Y_s \cdot d\gamma_s \text{ on } [t, t'], \text{ started at } A \in \overline{T(V)}$$

admits $Y_s = A \cdot S(\gamma)_{[t,s]}$ as unique solution, where the product \cdot on $\overline{T(V)}$ is defined for any $v = (v_0, v_1, \dots)$ and $w = (w_0, w_1, \dots)$ in $\overline{T(V)}$ as the element $v \cdot w = (z_0, z_1, \dots)$ in $\overline{T(V)}$ such that

$$z_k = \sum_{j=0}^k v_j \otimes w_{k-j} \in V^{\otimes k}, \quad \forall k \in \mathbb{N} \cup \{0\},$$

and where $d\gamma_t$ is identified with $(0, d\gamma_t, 0, \dots) \in \overline{T(V)}$. In particular, the signature $S(\gamma)_{t,t'}$ is the unique solution to (3.7) with initial condition $A = \mathbf{1} := (1, 0, 0, \dots) \in \overline{T(V)}$.

Proof. For any $a \in V$, the map $f_a : \overline{T(V)} \rightarrow L(V, \overline{T(V)})$ given by $f_a(Y) = Y \cdot a$ is well defined and linear. Thus, it can easily be seen to satisfy the conditions of the classical versions of Picard-Lindelöf theorem for Young CDEs (e.g. [36, Thm. 1.28]). Therefore, (3.7) admits a unique solution in $\overline{T(V)}$.

The fact that any solution $Y_s = (Y_s^0, Y_s^1, Y_s^2, \dots)$ to (3.7) in $\overline{T(V)}$ must satisfy $Y_s^k = (A \cdot S(\gamma)_{t,s})^k$ for any k in \mathbb{N} and any s in $[t, t']$ can be shown by a simple induction on k . \square

Remark 3.3. By construction, the iterated integrals in (3.6) admits the following recursive structure

$$S(\gamma)_{t,t'}^{(\omega)} = \int_t^{t'} S(\gamma)_{t,t'}^{(\omega_-)} d\gamma_u^{(i_m)}, \quad \text{where } \omega_- = (i_1, \dots, i_{m-1}).$$

Thus the coordinates of the signature when taken in isolation, do not satisfy a differential equation.

We can now define the signature kernel as the inner product of two signatures.

Definition 3.4. Let $0 \leq s < s' \leq T$, $0 \leq t < t' \leq T$. Define the signature kernel $\kappa_{sig} : \text{BV}([t, t'], V) \times \text{BV}([s, s'], V) \rightarrow \mathbb{R}$ as the following symmetric function

$$\kappa_{sig}(\gamma, \tau) = \langle S(\gamma)_{t,t'}, S(\tau)_{s,s'} \rangle_{\overline{T(V)}} = \sum_{k=0}^{\infty} \sum_{\omega \in \mathbb{W}: |\omega|=k} S(\gamma)_{t,t'}^{(\omega)} S(\tau)_{s,s'}^{(\omega)}.$$

Remark 3.5. Note that the two paths γ, τ need not be defined on the same time interval. As for the norms, we will omit the subscript t, t' when it is clear from the context.

It was proved in [44] that signature kernel can be realised as the solution of path-dependent PDE as stated in the next lemma; this effectively provides a *kernel trick* for the signature kernel, i.e. a way of computing it without reference to the signature map.

Lemma 3.6. [44, Theorem 2.5] Let $[s, s'], [t, t'] \in [0, T]$ and $\gamma \in \text{BV}([t, t'], V)$, $\tau \in \text{BV}([s, s'], V)$ be two paths. Then, the signature kernel of γ, τ is realised as the solution of the following two-parameter integral equation

$$(3.8) \quad \kappa_{sig}(\gamma, \tau) = \kappa_{sig}(\gamma, \tau)_{t',s'} \quad \text{where} \quad \kappa_{sig}(\gamma, \tau)_{r,q} = 1 + \int_s^q \int_t^r \kappa_{sig}(\gamma, \tau)_{u,v} \langle d\gamma_u, d\tau_v \rangle_V,$$

and with boundary conditions $\kappa_{sig}(\gamma, \tau)_{t,q} = \kappa_{sig}(\gamma, \tau)_{r,s} = 1$ for all $q \in [s, s'], r \in [t, t']$.

Several finite difference schemes are available for numerically evaluating solutions to equation (3.8), see [44, Section 3.1] for details.

3.3. Signature kernels from static kernels. In view of data science applications, instead of taking inner products of signatures directly, it might be beneficial to first lift paths in V to paths with values in the RKHS \mathcal{H}_g of a static kernel $g : V \times V \rightarrow \mathbb{R}$. More precisely, for any $\gamma \in \text{BV}([t, t'], V)$, with $0 \leq t < t' \leq T$, denote by $\gamma^g : [t, t'] \rightarrow \mathcal{H}_g$ the path defined for any $r \in [t, t']$ as $\gamma_r^g = g(\gamma_r, \cdot)$. For the signature of a lifted path γ^g to be well-defined we will require that the latter is at least continuous and of bounded variation, i.e. $\gamma^g \in \text{BV}([t, t'], \mathcal{H}_g)$. The following result gives a sufficient condition to ensure this and immediately follows from the definitions of bounded variation and Lipschitz-continuity. In particular, in the experimental section, we will make use of the *radial basis function* kernel as g , which is C^∞ and therefore Lipschitz-continuous.

Lemma 3.7. Let $g : V \times V \rightarrow \mathbb{R}$ be a kernel such that for any $v \in V$ the canonical feature map $g(v, \cdot)$ is Lipschitz-continuous and denote by \mathcal{H}_g its RKHS. Then, for any path $\gamma \in \text{BV}([t, t'], V)$, the lifted path $\gamma^g : [t, t'] \rightarrow \mathcal{H}_g$ defined for any $r \in [t, t']$ as $\gamma_r^g = g(\gamma_r, \cdot)$ is continuous and of bounded variation.

Once paths in $\text{BV}([t, t'], V)$ are lifted to paths in $\text{BV}([t, t'], \mathcal{H}_g)$ one can compute their signatures and take their inner products in the same ways as before. More precisely, given any path $\gamma \in \text{BV}([t, t'], V)$, the signature $S(\gamma^g)$ of the lifted path $\gamma^g \in \text{BV}([t, t'], \mathcal{H}_g)$ is a well-defined element of the closure of $T(\mathcal{H}_g)$. This yields the following definition of lifted signature kernel.

Definition 3.8. Let $g : V \times V \rightarrow \mathbb{R}$ be a kernel such that for any $v \in V$ the canonical feature map $g(v, \cdot)$ is Lipschitz-continuous and denote by \mathcal{H}_g its RKHS. Then, for any $[s, s'], [t, t'] \subset [0, T]$, the

g -lifted signature kernel $\kappa_{\text{sig}}^g : \text{BV}([t, t'], V) \times \text{BV}([s, s'], V) \rightarrow \mathbb{R}$ is defined as follows

$$\kappa_{\text{sig}}^g(\gamma, \tau) = \langle S(\gamma^g)_{t, t'}, S(\tau^g)_{s, s'} \rangle$$

where the inner product is taken in the closure of $T(\mathcal{H}_g)$.

Lemma 3.6 directly yields a kernel trick for the g -lifted signature kernel κ_{sig}^g :

$$(3.9) \quad \kappa_{\text{sig}}^g(\gamma, \tau) = \kappa_{\text{sig}}^g(\gamma, \tau)_{t', s'} \quad \text{where} \quad \kappa_{\text{sig}}^g(\gamma, \tau)_{r, q} = 1 + \int_s^q \int_t^r \kappa_{\text{sig}}^g(\gamma, \tau)_{u, v} \langle d\gamma_u^g, d\tau_v^g \rangle_{\mathcal{H}_g},$$

with the boundary conditions $\kappa_{\text{sig}}^g(\gamma, \tau)_{t, q} = \kappa_{\text{sig}}^g(\gamma, \tau)_{r, s} = 1$ for all $q \in [s, s']$, $r \in [t, t']$.

Remark 3.9. The integration with respect to \mathcal{H}_g paths is understood in the sense of

$$(3.10) \quad d\gamma_u^g = dg(\gamma_u, \cdot) = \sum_{i=1}^d \partial_{x^i} g(\gamma_u, \cdot) d\gamma_u^i.$$

3.4. Derivatives of static kernels. To approximate solutions to PPDEs, it will be necessary to differentiate a kernel with respect to its input variables. The RBF kernel g_σ , defined in (3.2), is clearly C^∞ and its n^{th} order derivative can be obtained using elementary calculus. For example, the first two derivatives, which are the only derivatives we will need to consider in the experimental section, admits the following expressions

$$(3.11) \quad \frac{\partial g_\sigma(x, y)}{\partial x} = \frac{y - x}{\sigma^2} \exp\left(-\frac{(x - y)^2}{2\sigma^2}\right)$$

$$(3.12) \quad \frac{\partial^2 g_\sigma(x, y)}{\partial x^2} = \frac{(x - y)^2 - \sigma^2}{\sigma^4} \exp\left(-\frac{(x - y)^2}{2\sigma^2}\right)$$

By the results in [55], for all $x \in V$, the derivatives of $g(x, \cdot)$ are elements of \mathcal{H}_g . Furthermore, [55, Theorem 1] states that, for any $h \in \mathcal{H}_g$ and any multi-index I ,

$$\langle \partial_{x^I}^{[I]} g(x, \cdot), h \rangle_{\mathcal{H}_g} = \partial_{x^I}^{[I]} h(x) = \partial_{x^I}^{[I]} \langle g(x, \cdot), h \rangle_{\mathcal{H}_g},$$

where $x^I := \{x^i, i \in I\}$. For $n, m \in \mathbb{N}$ let I, J be multi-indices on $\{1, \dots, n\}$ and $\{1, \dots, m\}$ respectively. In particular, we have

$$(3.13) \quad \langle \partial_{x^I}^n g(x, \cdot), \partial_{y^J}^m g(y, \cdot) \rangle_{\mathcal{H}_g} = \partial_{x^I}^n \partial_{y^J}^m \langle g(x, \cdot), g(y, \cdot) \rangle_{\mathcal{H}_g} = \partial_{x^I}^n \partial_{y^J}^m g(x, y).$$

We make the following standing assumption on the static kernel.

Assumption 3.10. $g : V \times V \rightarrow \mathbb{R}$ is four times differentiable and, for all $0 \leq n, m \leq 4$, we have

$$\overline{g_{nm}} := \sup_{x, y \in V} \left(\sum_{|I|=n, |J|=m} \left| \partial_{x^I}^n \partial_{y^J}^m g(x, y) \right|^2 \right)^{\frac{1}{2}} < \infty.$$

This assumption is satisfied for the Gaussian kernel (3.2).

3.5. Derivatives of signature kernels. In this section we will only consider paths on intervals of the form $[t, T]$, $[s, T]$, for $s, t \in \mathbb{T}$ and fix $V = \mathbb{R}^e$; this way the path spaces boil down to those introduced in Section 2.1. For two paths $\gamma \in \text{BV}_t$ and $\tau \in \text{BV}_s$, the directional derivative of the signature kernel with respect to its first variable and along the direction of a path $\eta \in \text{BV}_t$ is the Gateaux derivative (2.10) which perturbs the whole path:

$$(3.14) \quad \partial_\gamma^\eta \kappa_{\text{sig}}(\gamma, \tau) := \lim_{\varepsilon \rightarrow 0} \frac{\kappa_{\text{sig}}(\gamma_{\text{sig}} + \varepsilon \eta, \tau) - \kappa_{\text{sig}}(\gamma, \tau)}{\varepsilon}$$

It was shown in [45] that the directional derivative in equation (3.14) is such that

$$\partial_\gamma^\eta \kappa_{\text{sig}}(\gamma, \tau) = \partial_\gamma^\eta \kappa_{\text{sig}}(\gamma, \tau)_{T, T}$$

where the map $\partial_\gamma^\eta \kappa_{\text{sig}}(\gamma, \tau) : [t, T] \times [s, T] \rightarrow \mathbb{R}$ solves a two-parameter integral equation similar to (3.8)

$$(3.15) \quad \partial_\gamma^\eta \kappa_{\text{sig}}(\gamma, \tau)_{s', t'} = \int_s^{s'} \int_t^{t'} \partial_\gamma^\eta \kappa_{\text{sig}}(\gamma, \tau)_{u, v} \langle d\gamma_u, d\tau_v \rangle + \kappa_{\text{sig}}(\gamma, \tau)_{u, v} \langle d\eta_u, d\tau_v \rangle.$$

Leveraging (3.9), we can extend this PDE formulation to the lifted kernel and to its second derivative, as we show in Proposition 3.11. For all $\gamma \in \text{BV}_t$, define the map $G : \text{BV}_t \rightarrow \text{BV}([t, T], \mathcal{H}_g)$ as

$$G(\gamma)_u = \gamma_u^g = g(\gamma_u, \cdot), \quad \text{for all } u \in [t, T].$$

The proofs of all the results of this section are postponed to Appendix A.2 to ease the flow of the paper.

Proposition 3.11. *Let Assumption 3.10 holds and $\gamma \in \text{BV}_t, \tau \in \text{BV}_s$ for $s, t \in \mathbb{T}$. Then the following hold for all $s' \in [s, T]$ and $t' \in [t, T]$.*

(1) *For any $\eta \in \text{BV}_t$, the directional derivative of the lifted kernel satisfies the following PDE*

$$(3.16) \quad \partial_\gamma^\eta \kappa_{\text{sig}}^g(\gamma, \tau)_{s', t'} = \int_s^{s'} \int_t^{t'} \left(\partial_\gamma^\eta \kappa_{\text{sig}}^g(\gamma, \tau)_{u, v} \langle d\gamma_u^g, d\tau_v^g \rangle_{\mathcal{H}_g} + \kappa_{\text{sig}}^g(\gamma, \tau)_{u, v} \langle d\partial_\gamma^\eta G(\gamma)_u, d\tau_v^g \rangle_{\mathcal{H}_g} \right).$$

(2) *For any $\eta, \bar{\eta} \in \text{BV}_t$, the second derivative of the lifted kernel satisfies the following PDE*

$$(3.17) \quad \begin{aligned} \partial_\gamma^{\eta \bar{\eta}} \kappa_{\text{sig}}^g(\gamma, \tau)_{s', t'} &= \int_s^{s'} \int_t^{t'} \partial_\gamma^{\eta \bar{\eta}} \kappa_{\text{sig}}^g(\gamma, \tau)_{u, v} \langle d\gamma_u^g, d\tau_v^g \rangle_{\mathcal{H}_g} + \int_s^{s'} \int_t^{t'} \partial_\gamma^{\eta \bar{\eta}} \kappa_{\text{sig}}^g(\gamma, \tau)_{u, v} \langle d\partial_\gamma^{\eta \bar{\eta}} G(\gamma)_u, d\tau_v^g \rangle_{\mathcal{H}_g} \\ &\quad + \int_s^{s'} \int_t^{t'} \partial_\gamma^{\bar{\eta}} \kappa_{\text{sig}}^g(\gamma, \tau)_{u, v} \langle d\partial_\gamma^\eta G(\gamma)_u, d\tau_v^g \rangle_{\mathcal{H}_g} + \int_s^{s'} \int_t^{t'} \kappa_{\text{sig}}^g(\gamma, \tau)_{u, v} \langle d\partial_\gamma^{\eta \bar{\eta}} G(\gamma)_u, d\tau_v^g \rangle_{\mathcal{H}_g}. \end{aligned}$$

From the PDEs of Proposition 3.11 arise uniform estimates in $s, t \in \mathbb{T}$ for the signature kernel and its derivatives with respect to the 1-variation norms of $\gamma, \tau, \eta, \bar{\eta}$. The precise statements and bounds are gathered in Lemma A.2 in the Appendix. In Section 4 we will restrict the state space to a compact subspace of BV_0 with bounded 1-variation norm on which these estimates are uniformly bounded.

Remark 3.12. Expliciting the partial derivatives in (3.16) yields the equivalent formulation:

$$(3.18) \quad \begin{aligned} \partial_\gamma^\eta \kappa_{\text{sig}}^g(\gamma, \tau)_{s', t'} &= \int_s^{s'} \int_t^{t'} \partial_\gamma^\eta \kappa_{\text{sig}}^g(\gamma, \tau)_{u, v} \sum_{i, k=1}^e \partial_{x^i y^k}^2 g(\gamma_u, \tau_v) d\gamma_u^i d\tau_v^k \\ &\quad + \int_s^{s'} \int_t^{t'} \kappa_{\text{sig}}^g(\gamma, \tau)_{u, v} \sum_{i, j, k=1}^e \partial_{x^i y^j y^k}^3 g(\gamma_u, \tau_v) \eta_u^j d\gamma_u^i d\tau_v^k \\ &\quad + \int_s^{s'} \int_t^{t'} \kappa_{\text{sig}}^g(\gamma, \tau)_{u, v} \sum_{i, k=1}^e \partial_{x^i y^k}^2 g(\gamma_u, \tau_v) d\eta_u^i d\tau_v^k. \end{aligned}$$

Equation (3.16) is more conducive for numerical computations while Equation (3.18) will be used for the theoretical analysis. One can perform the same expansion for the second derivative, as can be seen in the proof (Equation (3.18)).

Remark 3.13. By iterating the same arguments one could prove that the pathwise derivative of the kernel at any order satisfies a linear PDE of the same kind. We do not pursue this here, but note in particular that $\kappa(\cdot, \tau) \in C^2$ for all $\tau \in \text{BV}(\mathbb{T}, \mathbb{R}^e)$.

Similarly, given two paths $\gamma, \eta \in \text{BV}_t$, we can define the directional derivative

$$(3.19) \quad \partial_\gamma^\eta S \circ G(\gamma)_{t, t'} = \partial_\gamma^\eta S(\gamma^g)_{t, t'} := \lim_{\varepsilon \rightarrow 0} \frac{S((\gamma + \varepsilon \eta)^g)_{t, t'} - S(\gamma^g)_{t, t'}}{\varepsilon},$$

as well as the second derivative with respect to another direction $\bar{\eta} \in \text{BV}(\mathbb{T}, V)$, denoted $\partial_{\gamma}^{\eta\bar{\eta}} S \circ G$. The following result is a corollary of Lemma A.1, which also provides estimates of the signature and its derivatives with respect to the 1-variation norms of $\gamma, \eta, \bar{\eta}$.

Proposition 3.14. *For all $\gamma, \eta, \bar{\eta} \in \text{BV}(\mathbb{T}, \mathbb{R}^e)$, $e \in \mathbb{N}$, the signature and its first two directional derivatives $\partial_{\gamma}^{\eta} S \circ G$ and $\partial_{\gamma}^{\eta\bar{\eta}} S \circ G$ exist and are continuous with respect to the C^0 -norm. Moreover, the derivatives satisfy the CDEs (A.9) and (A.12).*

This result is analogous to [25, Theorem 4.4] albeit in \mathcal{H}_g valued paths.

3.6. Computations of derivatives. In practice, signature kernels and their derivatives can be computed with a single call to a PDE solver. To see this, for any $\gamma, \tau, \eta, \bar{\eta} \in \text{BV}$ and any $s, t \in \mathbb{T}$ set

$$\mathbb{A}_{s,t}(\gamma, \tau) = \langle \partial_t \gamma_s^g, \partial_t \tau_t^g \rangle_{\mathcal{H}_g}, \quad \mathbb{A}_{s,t}^{\eta}(\gamma, \tau) = \langle \partial_t \partial_{\gamma}^{\eta} G(\gamma)_s, \partial_t \tau_t^g \rangle_{\mathcal{H}_g}, \quad \mathbb{A}_{s,t}^{\eta\bar{\eta}}(\gamma, \tau) = \langle \partial_t \partial_{\gamma}^{\eta\bar{\eta}} G(\gamma)_s, \partial_t \tau_t^g \rangle_{\mathcal{H}_g}$$

Then, the augmented variable

$$\mathbf{K}_{s,t}^{\eta\bar{\eta}}(\gamma, \tau) = \begin{pmatrix} \kappa_{\text{sig}}^g(\gamma, \tau)_{s,t} \\ \partial_{\gamma}^{\eta} \kappa_{\text{sig}}^g(\gamma, \tau)_{s,t} \\ \partial_{\gamma}^{\bar{\eta}} \kappa_{\text{sig}}^g(\gamma, \tau)_{s,t} \\ \partial_{\gamma}^{\eta\bar{\eta}} \kappa_{\text{sig}}^g(\gamma, \tau)_{s,t} \end{pmatrix} \in \mathbb{R}^4$$

satisfies the following linear system of hyperbolic PDEs

$$(3.20) \quad \frac{\partial^2}{\partial_s \partial_t} \mathbf{K}_{s,t}^{\eta\bar{\eta}}(\gamma, \tau) = \begin{pmatrix} \mathbb{A}_{s,t}(\gamma, \tau) & 0 & 0 & 0 \\ \mathbb{A}_{s,t}^{\eta}(\gamma, \tau) & \mathbb{A}_{s,t}(\gamma, \tau) & 0 & 0 \\ \mathbb{A}_{s,t}^{\bar{\eta}}(\gamma, \tau) & 0 & \mathbb{A}_{s,t}(\gamma, \tau) & 0 \\ \mathbb{A}_{s,t}^{\eta\bar{\eta}}(\gamma, \tau) & \mathbb{A}_{s,t}^{\eta}(\gamma, \tau) & \mathbb{A}_{s,t}^{\bar{\eta}}(\gamma, \tau) & \mathbb{A}_{s,t}(\gamma, \tau) \end{pmatrix} \mathbf{K}_{s,t}^{\eta\bar{\eta}}(\gamma, \tau),$$

with boundary conditions

$$\mathbf{K}_{0,t}^{\eta\bar{\eta}}(\gamma, \tau) = \mathbf{K}_{s,0}^{\eta\bar{\eta}}(\gamma, \tau) = (1, 0, 0, 0)^{\top}.$$

Remark 3.15. In the case where $g = \text{id}$ the system of PDEs (3.20) simplifies as $\mathbb{A}_{s,t}^{\eta}(\gamma, \tau) = \mathbb{A}_{s,t}(\eta, \tau)$ and $\mathbb{A}_{s,t}^{\eta\bar{\eta}} = 0$. For more general static kernels g , the derivative \mathbb{A} might not be available in close form. In that case we can approximate them by finite difference as follows

$$\mathbb{A}_{s,t}(\gamma, \tau) \approx \frac{1}{\Delta s \Delta t} \left(g(\gamma_s, \tau_t) - g(\gamma_s, \tau_{t-\Delta t}) - g(\gamma_{s-\Delta s}, \tau_t) + g(\gamma_{s-\Delta s}, \tau_{t-\Delta t}) \right).$$

Similarly for \mathbb{A}^{η} and $\mathbb{A}^{\eta, \bar{\eta}}$.

3.7. Universal product kernels. The solutions to our PPDEs are defined on the product space $\Omega = \mathbb{T} \times \mathbb{R}^d \times \text{BV}$. Therefore, we will need to consider product kernels indexed on such product spaces. The following result provides a simple way of constructing kernels that are universal on a product space as products of kernel on the individual factor spaces.

Lemma 3.16. [9, Lemma 5.2] *Let κ_1, κ_2 be two cc-universal kernels indexed on Ω_1, Ω_2 respectively. Then the product kernel defined for any $\omega_1, \omega'_1 \in \Omega_1$ and $\omega_2, \omega'_2 \in \Omega_2$ as*

$$\kappa((\omega_1, \omega_2), (\omega'_1, \omega'_2)) := \kappa_1(\omega_1, \omega'_1) \kappa_2(\omega_2, \omega'_2)$$

is cc-universal on $\Omega_1 \times \Omega_2$.

Recall that the domains Ω of our PPDEs are of the form $\Omega = \mathbb{T} \times \mathbb{R}^d \times \text{BV}_0 \subset \mathbb{R}^{d+1} \times \text{BV}_0$. Thus, it would be tempting to consider product kernels of the form $\tilde{\kappa}^{k,g} : \Omega \times \Omega \rightarrow \mathbb{R}$ defined for any $(s, x, \gamma), (t, y, \tau) \in \Omega$ as

$$\tilde{\kappa}^{h,g}((s, x, \gamma), (t, y, \tau)) := k((s, x), (t, y)) \kappa_{\text{sig}}^g(\gamma, \tau),$$

where $k : \mathbb{R}^{d+1} \times \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ is some cc-universal kernel on \mathbb{R}^{d+1} such as the RBF kernel in (3.2), and κ_{sig}^g is a g -lifted signature kernel.

However, signature kernels are only universal when restricted to a class of paths where the signature is an injective map, as we explain next.

Let $BV_0^0 \subset BV_0$ be the subset of continuous bounded variation paths that are time-augmented and started at the origin 0, endowed with the 1-variation topology. Denote by $\mathcal{H}_{\kappa_{\text{sig}}^g}$ the RKHS associated to the signature kernel κ_{sig}^g . The next result states that, when restricted to a compact subset $B \subset BV_0^0$, elements of $\mathcal{H}_{\kappa}|_B$ are dense in $\mathcal{C}(B)$ with the topology of uniform convergence. In other words, κ_{sig}^g is cc-universal on BV_0^0 .

Lemma 3.17. *Let $B \subset BV_0^0$ be compact in 1-variation. Then, $\mathcal{H}_{\kappa}|_B$ is a dense subset of $\mathcal{C}(B)$ with the topology of uniform convergence.*

The proof of this statement is analogous to [11, Proposition 3.3] with the difference that the point-separation assumption in the Stone-Weierstrass theorem is justified by noting that $S(\gamma) = S(\tau) \iff \gamma = \tau$ when $\gamma, \tau \in BV_0^0$.

It is important to note that the above result no longer holds if one considers compact subsets of BV_0 that are not time-augmented and that do not share a common origin. This is due the fact that the signature map is no longer injective on BV_0 , thus the point-separation assumption used in the proof of Lemma 3.17 invoking the Stone-Weierstrass Theorem no longer holds. However, in this paper we are interested in paths that might not have a common starting point. To remedy this issue, we modify the signature kernel κ_{sig}^g and define the kernel $\kappa_{\text{sig}}^{g,\ell} : BV_s \times BV_t \rightarrow \mathbb{R}$ as

$$\kappa_{\text{sig}}^{g,\ell}(\gamma, \tau) = \ell(\gamma_s, \tau_t) \kappa_{\text{sig}}^g(\gamma - \gamma_s, \tau - \tau_t)$$

where $\ell : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is a cc-universal kernel. We assume that the associated feature map $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$ is bounded, Lipschitz continuous, twice differentiable and with bounded, Lipschitz continuous derivatives. By Lemma 3.17 κ_{sig}^g is cc-universal on BV_0 . Thus, by Lemma 3.16 $\kappa_{\text{sig}}^{g,\ell}$ is cc-universal on BV .

Hence, we define the product kernel $\kappa : \Omega \times \Omega \rightarrow \mathbb{R}$ for any $(s, x, \gamma), (t, y, \tau) \in \Omega$ as

$$(3.21) \quad \kappa((s, x, \gamma), (t, y, \tau)) = \kappa^{k,g,\ell}((s, x, \gamma), (t, y, \tau)) = k((s, x), (t, y)) \kappa_{\text{sig}}^{g,\ell}(\gamma|_{[s,T]}, \tau|_{[t,T]}),$$

where $k : \mathbb{R}^{d+1} \times \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ is a cc-universal kernel on \mathbb{R}^{d+1} . We assume that the associated feature map is bounded, Lipschitz continuous, twice differentiable and with bounded, Lipschitz continuous derivatives.

We note that, for technical reasons, we have to work on a compact subspace of BV_0 with respect to the $1 < p$ -variation norm and on which the 1-variation norm is bounded. This results in all the estimates of Lemmas A.2 and A.1 to be bounded and enables the use of Arzelà-Ascoli theorem. As we recalled in the introduction, for all $R > 0$ and $\alpha > 0$, the set $B(R) := \{\gamma \in BV_0 : \|\gamma\|_{1-var,\alpha} \leq R\}$ is compact in BV_0 with respect to the p -variation topology, for all $p > 1$. Compact sets of Ω with respect to \mathbf{d} are thus typically of the form $\mathfrak{K} = \mathbb{T} \times \prod_{i=1}^d [a_i, b_i] \times B(R)$, for $a_i, b_i \in \mathbb{R}$, $i = 1, \dots, d$ and $R > 0$.

We denote by \mathcal{H} the RKHS of the kernel $\kappa : \mathfrak{K} \times \mathfrak{K} \rightarrow \mathbb{R}$ restricted to act on \mathfrak{K} . We define the feature map $\Phi : \mathfrak{K} \rightarrow \mathcal{H}$ as $\Phi(\omega) = \kappa(\omega, \cdot)$. Since the two kernels h and $\kappa_{\text{sig}}^{g,\ell}$ are continuously twice differentiable, as we have seen in the previous section, so is their product. This implies that for each $\omega \in \mathfrak{K}$, the map $\kappa(\omega, \cdot)$ belongs to $C^{\infty,\infty,2}(\mathfrak{K})$.

Moreover, as a consequence of Lemmas 3.16 and 3.17, the product kernel κ is cc-universal on $\mathfrak{K} \times \mathfrak{K}$. In other words, \mathcal{H} is dense in $\mathcal{C}(\mathfrak{K})$ with respect to the topology of uniform convergence.

4. MAIN THEORETICAL RESULTS

This section aims at answering the questions posed in Section 2.5 and restricts our domain of study to pathwise derivatives of order two or lower. This framework is relevant for the backward Kolmogorov equations we presented in Section 2.3 and can rely on the results shown in Proposition 3.11. More

precisely we fix $\bar{N} \in \mathbb{N}$ and consider the operator \mathcal{L}_t in (2.4) with

$$\mathcal{N} := \{n = (n_1, n_2, n_3) \in \mathbb{N}_0^3 : n_1 \leq \bar{N}, n_2 \leq \bar{N}, n_3 \leq 2\}.$$

4.1. Optimal recovery and well-posedness. This section aims at answering the questions (1)-(3) of Section 2.5 regarding the well-posedness of the optimisation problem (2.13). Our first result verifies that elements of \mathcal{H} are indeed twice differentiable, and therefore belong to the domain of the PPDE operator \mathcal{L}_t . This answers Question (1) in 2.5. The proof builds on the estimates derived through blood and sweat in Lemma A.2. The technical hurdle consists in interchanging the infinite series and the derivative operator. As advertised earlier, this operation requires the restriction of \mathcal{H} to a compact subset of Ω called \mathfrak{K} . Even though \mathfrak{K} is closed, the time derivative of $h \in \mathcal{H}$ is well-defined since it only perturbs to the right and we set $\partial_t u(T, x, \gamma) = 0$. Moreover the spatial and pathwise derivatives are also well-defined over \mathfrak{K} since the kernel κ is defined over $\Omega \times \Omega$.

Proposition 4.1. *Let Assumption 3.10 hold for g . For all $h \in \mathcal{H}$ such that $h = \sum_{i=1}^{\infty} \alpha_i \kappa(\omega_i, \cdot)$, with $(\alpha_i)_{i \in \mathbb{N}}$ in \mathbb{R} and $(\omega_i)_{i \in \mathbb{N}}$ in \mathfrak{K} , for all $\partial \in \{\partial_t^{n_1} \partial_x^{n_2} \partial_\gamma^\eta : (n_1, n_2, n_3) \in \mathcal{N}, \eta \in \text{BV}_0^{\otimes n_3}\}$ and $\omega = (t, x, \gamma) \in \mathfrak{K}$, we have*

$$\partial h(\omega) = \sum_{i=1}^{\infty} \alpha_i \partial \kappa(\omega_i, \omega).$$

In particular the domain of \mathcal{L}_t is included in \mathcal{H} for all $t \in [0, T]$.

Proof. Let $h \in \mathcal{H}$ such that $h = \sum_{i=1}^{\infty} \alpha_i \kappa(\omega_i, \cdot)$, with $(\alpha_i)_{i \in \mathbb{N}}$ in \mathbb{R} and $(\omega_i)_{i \in \mathbb{N}}$ in \mathfrak{K} . We prove the claim only for the pathwise derivative $\partial \in \{\partial_\gamma^\eta, \widehat{\partial}_\gamma^\eta\}$, with $\eta, \bar{\eta} \in \text{BV}_0$, and note that the other cases can be proved in an identical fashion, as the Gaussian kernel k is infinitely many times differentiable and its derivatives are uniformly bounded. An important observation is that for all $(s, x, \gamma), (t, y, \tau) \in \mathfrak{K}$ we have

$$\widehat{\partial}_\gamma^\eta \kappa((s, x, \gamma), (t, y, \tau)) = k((s, x), (t, y)) \partial_\gamma^{\eta|_{[t, T]}} \kappa_{\text{sig}}^{g, \ell}(\gamma|_{[t, T]}, \tau).$$

This highlights the link between the two types of pathwise derivatives $\widehat{\partial}_\gamma^\eta$ and ∂_γ^η .

(∂_γ^η) For all $n \in \mathbb{N}$, let $h_n := \sum_{i=1}^n \alpha_i \kappa(\omega_i, \cdot)$ hence h_n tends to f pointwise and in RKHS norm as $n \rightarrow \infty$. For all $\omega \in \mathfrak{K}$ and $\eta \in \text{BV}(\mathbb{T}, \mathbb{R}^e)$, we have and we define

$$\partial_\gamma^\eta h_n(\omega) = \sum_{i=1}^n \alpha_i \partial_\gamma^\eta \kappa(\omega_i, \omega), \quad h'(\omega) := \sum_{i=1}^{\infty} \alpha_i \partial_\gamma^\eta \kappa(\omega_i, \omega),$$

and we will prove that $\partial_\gamma^\eta h_n$ converges to h' uniformly on \mathfrak{K} . Let $n \in \mathbb{N}$, the Cauchy-Schwarz inequality allows us to disentangle the two limits:

(4.1)

$$\begin{aligned} |\partial_\gamma^\eta h_n(\omega) - h'(\omega)| &= \left| \sum_{i=n+1}^{\infty} \alpha_i \partial_\gamma^\eta \kappa(\omega_i, \omega) \right| = \left| \sum_{i=n+1}^{\infty} \alpha_i \lim_{\varepsilon \rightarrow 0} \frac{\kappa(\omega_i, (t, x, \gamma + \varepsilon \eta)) - \kappa(\omega_i, (t, x, \gamma))}{\varepsilon} \right| \\ &= \lim_{N \rightarrow \infty} \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \left| \sum_{i=n+1}^N \alpha_i \left\langle \Phi(\omega_i), \Phi(t, x, \gamma + \varepsilon \eta) - \Phi(t, x, \gamma) \right\rangle \right| \\ &= \lim_{N \rightarrow \infty} \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \left| \left\langle \sum_{i=n+1}^N \alpha_i \Phi(\omega_i), \Phi(t, x, \gamma + \varepsilon \eta) - \Phi(t, x, \gamma) \right\rangle \right| \\ &\leq \lim_{N \rightarrow \infty} \left\| \sum_{i=n+1}^N \alpha_i \Phi(\omega_i) \right\|_{\mathcal{H}} \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \|\Phi(t, x, \gamma + \varepsilon \eta) - \Phi(t, x, \gamma)\|_{\mathcal{H}}. \end{aligned}$$

We observe that

$$\begin{aligned}
\left\| \sum_{i=n+1}^{\infty} \alpha_i \Phi(\omega_i) \right\|_{\mathcal{H}}^2 &= \left\langle \sum_{i=n+1}^{\infty} \alpha_i \Phi(\omega_i), \sum_{j=n+1}^{\infty} \alpha_j \Phi(\omega_j) \right\rangle_{\mathcal{H}} = \sum_{i=n+1}^{\infty} \alpha_i \sum_{j=n+1}^{\infty} \alpha_j \langle \Phi(\omega_i), \Phi(\omega_j) \rangle_{\mathcal{H}} \\
&= \sum_{i,j=n+1}^{\infty} \alpha_i \alpha_j \kappa(\omega_i, \omega_j) \\
&= \|h - h_n\|_{\mathcal{H}}^2,
\end{aligned}$$

which tends to zero. Furthermore,

$$\begin{aligned}
&\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon^2} \|\Phi(t, x, \gamma + \varepsilon \eta) - \Phi(t, x, \gamma)\|_{\mathcal{H}}^2 \\
&= k((t, x), (t, x)) \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon^2} \left(\kappa_{\text{sig}}^{g, \ell}(\gamma + \varepsilon \eta, \gamma + \varepsilon \eta) - 2\kappa_{\text{sig}}^{g, \ell}(\gamma + \varepsilon \eta, \gamma) + \kappa_{\text{sig}}^{g, \ell}(\gamma, \gamma) \right) \\
&= k((t, x), (t, x)) \left(\ell(\gamma_t, \gamma_t) \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon^2} \|S((\gamma + \varepsilon \eta)^g) - S(\gamma^g)\|_{\overline{T(\mathcal{H}_g)}}^2 \right. \\
&\quad \left. + \kappa_{\text{sig}}^g(\gamma - \gamma_t, \gamma - \gamma_t) \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon^2} |\varphi(\gamma_t + \varepsilon \eta_t) - \varphi(\gamma_t)|^2 \right) \\
&= k((t, x), (t, x)) \left(\ell(\gamma_t, \gamma_t) \|\partial_{\gamma}^{\eta} S(\gamma^g)\|_{\overline{T(\mathcal{H}_g)}}^2 + \kappa_{\text{sig}}^g(\gamma - \gamma_t, \gamma - \gamma_t) |\partial_{\gamma_t}^{\eta_t} \varphi(\gamma_t)|^2 \right)
\end{aligned}$$

which is uniformly bounded over all $\omega \in \mathfrak{K}$ by Lemma A.1. We used dominated convergence in the last line which follows from (A.2). This entails the convergence of $\partial_{\gamma}^{\eta} h_n$ towards h' is uniform in \mathfrak{K} ; therefore h is differentiable and $\partial_{\gamma}^{\eta} h = h' = \sum_{i=1}^{\infty} \alpha_i \partial_{\gamma}^{\eta} \kappa(\omega_i, \cdot)$.

$(\partial_{\gamma}^{\eta \bar{\eta}})$ For all $\omega \in \mathfrak{K}$ and $\eta, \bar{\eta} \in \text{BV}(\mathbb{T}, \mathbb{R}^e)$, we have

$$\partial_{\gamma}^{\eta \bar{\eta}} h_n(\omega) := \sum_{i=1}^n \alpha_i \partial_{\gamma}^{\eta \bar{\eta}} \kappa(\omega_i, \omega), \quad h''(\omega) := \sum_{i=1}^{\infty} \alpha_i \partial_{\gamma}^{\eta \bar{\eta}} \kappa(\omega_i, \omega),$$

and we will prove that $\lim_{n \rightarrow \infty} \partial_{\gamma}^{\eta \bar{\eta}} h_n = h''$. First note that

$$\begin{aligned}
\partial_{\gamma}^{\eta \bar{\eta}} \kappa(\omega_i, \omega) &= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \left(\partial_{\gamma}^{\bar{\eta}} \kappa(\omega_i, (t, x, \gamma + \varepsilon \eta)) - \partial_{\gamma}^{\bar{\eta}} \kappa(\omega_i, (t, x, \gamma)) \right) \\
&= \lim_{\varepsilon \rightarrow 0} \lim_{\delta \rightarrow 0} \frac{1}{\varepsilon \delta} \left(\kappa(\omega_i, (t, x, \gamma + \varepsilon \eta + \delta \bar{\eta})) - \kappa(\omega_i, (t, x, \gamma + \varepsilon \eta)) - \kappa(\omega_i, (t, x, \gamma + \delta \bar{\eta})) + \kappa(\omega_i, (t, x, \gamma)) \right) \\
&= \lim_{\varepsilon \rightarrow 0} \lim_{\delta \rightarrow 0} \frac{1}{\varepsilon \delta} \left\langle \Phi(\omega_i), \Phi(t, x, \gamma + \varepsilon \eta + \delta \bar{\eta}) - \Phi(t, x, \gamma + \varepsilon \eta) - \Phi(t, x, \gamma + \delta \bar{\eta}) + \Phi(t, x, \gamma) \right\rangle.
\end{aligned}$$

Applying the same approach as for the first derivative, we obtain by Cauchy-Schwarz inequality

$$\begin{aligned}
|\partial_{\gamma}^{\eta \bar{\eta}} h_n(\omega) - h''(\omega)| &= \left| \sum_{i=n+1}^{\infty} \alpha_i \partial_{\gamma}^{\eta \bar{\eta}} \kappa(\omega_i, \omega) \right| \\
&= \lim_{N \rightarrow \infty} \lim_{\varepsilon \rightarrow 0} \lim_{\delta \rightarrow 0} \frac{1}{\varepsilon \delta} \left| \left\langle \sum_{i=n+1}^N \alpha_i \Phi(\omega_i), \Phi(t, x, \gamma + \varepsilon \eta + \delta \bar{\eta}) - \Phi(t, x, \gamma + \varepsilon \eta) - \Phi(t, x, \gamma + \delta \bar{\eta}) + \Phi(t, x, \gamma) \right\rangle \right| \\
&\leq \lim_{N \rightarrow \infty} \|h - h_n\|_{\mathcal{H}} \lim_{\varepsilon \rightarrow 0} \lim_{\delta \rightarrow 0} \frac{1}{\varepsilon \delta} \|\Phi(t, x, \gamma + \varepsilon \eta + \delta \bar{\eta}) - \Phi(t, x, \gamma + \varepsilon \eta) - \Phi(t, x, \gamma + \delta \bar{\eta}) + \Phi(t, x, \gamma)\|_{\mathcal{H}} \\
&= \lim_{N \rightarrow \infty} \|h - h_n\|_{\mathcal{H}} k((t, x), (t, x)) \left(|\varphi(\gamma_t)| \|\partial_{\gamma}^{\eta \bar{\eta}} S(\gamma^g)\|_{\overline{T(\mathcal{H}_g)}} + |\partial_{\gamma_t}^{\eta_t} \varphi(\gamma_t)| \|\partial_{\gamma}^{\bar{\eta}} S(\gamma^g)\|_{\overline{T(\mathcal{H}_g)}} \right. \\
&\quad \left. + |\partial_{\gamma_t}^{\bar{\eta}_t} \varphi(\gamma_t)| \|\partial_{\gamma}^{\eta} S(\gamma^g)\|_{\overline{T(\mathcal{H}_g)}} + |\partial_{\gamma_t}^{\eta_t \bar{\eta}_t} \varphi(\gamma_t)| \|S(\gamma^g)\|_{\overline{T(\mathcal{H}_g)}} \right),
\end{aligned}$$

where we used dominated convergence as for the first derivative. Lemma A.1 thus ensures that $\partial_\gamma^{\eta\bar{\eta}} h_n$ converges towards h'' uniformly in \mathfrak{K} and thus h is twice differentiable with $\partial_\gamma^{\eta\bar{\eta}} h = h'' = \sum_{i=1}^\infty \alpha_i \partial_\gamma^{\eta\bar{\eta}} \kappa(\omega_i, \cdot)$. \square

Now that we have checked that the constraints of the optimal recovery problem (2.13) make sense, we arrive at the following representer theorem, which gives an answer to Question (2) of our list 2.5.

Theorem 4.2. *Let Assumption 3.10 hold for g . For all $n, m \in \mathbb{N}$ and collocation points $\omega_1, \dots, \omega_{m+n} \in \mathfrak{K}$, there exist weights $\alpha_1, \dots, \alpha_{m+n} \in \mathbb{R}$ such that the optimal recovery problem*

$$(4.2) \quad \min_{u \in \mathcal{H}} \frac{1}{2} \|u\|_{\mathcal{H}}^2 \quad \text{s.t.} \quad \begin{cases} \mathcal{L}_t u(\omega^i) = f(\omega^i) & i = 1, \dots, m \\ u(\omega^i) = \phi(\omega^i) & i = m+1, \dots, m+n. \end{cases}$$

associated to the RKHS \mathcal{H} is well-posed and has a unique minimiser $u_{m,n} \in \mathcal{H}$ of the form

$$(4.3) \quad u_{m,n} = \sum_{i=1}^{m+n} \alpha_i \kappa(\omega^i, \cdot).$$

Proof. The proof is similar to the proof of classical representer theorems with the addition of pointwise observation of higher order derivatives. The existence and uniqueness of the minimiser of the optimal recovery problem (4.2) is guaranteed by the fact that the Lagrangian functional

$$(4.4) \quad u \mapsto \frac{1}{2} \|u\|_{\mathcal{H}}^2 + \sum_{i=1}^m \lambda_i \mathcal{L}_t u(\omega^i) + \sum_{i=m+1}^{m+n} \lambda_i (u(\omega^i) - \phi(\omega^i))^2, \quad \lambda_1, \dots, \lambda_{m+n} > 0$$

is continuous and convex from \mathcal{H} to \mathbb{R} . Let $\mathcal{A} := \text{Span}\{\kappa(\omega^i, \cdot) : i = 1, \dots, m+n\}$. Because $\mathcal{H} = \mathcal{A} \oplus \mathcal{A}^\perp$, any $u \in \mathcal{H}$ admits the following decomposition $u = \hat{u} + u^\perp$, where $\hat{u} = \sum_{i=1}^{m+n} \alpha_i \kappa(\omega^i, \cdot) \in \mathcal{A}$ for some scalars $\alpha_1, \dots, \alpha_{m+n} \in \mathbb{R}$, and $u^\perp \in \mathcal{A}^\perp$. By the reproducing property of κ , for any $j = 1, \dots, m+n$ the following relation holds

$$u(\omega^j) = \langle u, \kappa(\omega^j, \cdot) \rangle_{\mathcal{H}} = \sum_{i=1}^{m+n} \alpha_i \kappa(\omega^i, \omega^j) + \langle u^\perp, \kappa(\omega^j, \cdot) \rangle_{\mathcal{H}} = \hat{u}(\omega^j),$$

where the last equality follows from the fact that $u^\perp \in \mathcal{A}^\perp$. Hence, the following equality holds

$$\sum_{i=1}^m \lambda_i \mathcal{L}_t u(\omega^i) + \sum_{i=m+1}^{m+n} \lambda_i (u(\omega^i) - \phi(\omega^i))^2 = \sum_{i=1}^m \lambda_i \mathcal{L}_t \hat{u}(\omega^i) + \sum_{i=m+1}^{m+n} \lambda_i (\hat{u}(\omega^i) - \phi(\omega^i))^2.$$

Since for any $u \in \mathcal{H}$ the functional $u \mapsto \frac{1}{2} \|u\|_{\mathcal{H}}^2$ is strictly monotonically increasing, one has $\frac{1}{2} \|u\|_{\mathcal{H}}^2 \geq \frac{1}{2} \|\hat{u}\|_{\mathcal{H}}^2$, which leads to the required inequality

$$\frac{1}{2} \|\hat{u}\|_{\mathcal{H}}^2 + \sum_{i=1}^m \lambda_i \mathcal{L}_t \hat{u}(\omega^i) + \sum_{i=m+1}^{m+n} \lambda_i (\hat{u}(\omega^i) - \phi(\omega^i))^2 \leq \frac{1}{2} \|u\|_{\mathcal{H}}^2 + \sum_{i=1}^m \lambda_i \mathcal{L}_t u(\omega^i) + \sum_{i=m+1}^{m+n} \lambda_i (u(\omega^i) - \phi(\omega^i))^2.$$

\square

Remark 4.3. Note that $(\alpha_i)_{i=1}^{m+n}$ crucially depend on the collocation points and on m and n ; in particular they have no reason to remain stable for different values of m, n . This makes comparison among $(u_{m,n})$ much trickier. As another avenue for future research, we note that orthogonalising the collocation points in a way that the $(\alpha_i)_i$ remain constant for all m, n could lead to better control of the norm.

Note that because the functional $u_{m,n}$ defined in (4.3) is an element of the RKHS \mathcal{H} , by the reproducing property we have the following expression of the squared norm

$$(4.5) \quad \|u_{m,n}\|_{\mathcal{H}}^2 = \sum_{i=1}^{m+n} \sum_{j=1}^{m+n} \alpha_i \kappa(\omega^i, \omega^j) \alpha_j = \boldsymbol{\alpha}^\top \mathcal{K} \boldsymbol{\alpha},$$

where $\alpha = (\alpha_1, \dots, \alpha_{m+n}) \in \mathbb{R}^{m+n}$ and $\mathcal{K} \in \mathbb{R}^{(m+n) \times (m+n)}$ is the matrix with entries $\mathcal{K}_{i,j} = \kappa(\omega^i, \omega^j)$. Consequently, plugging the expression in (4.5) into the original optimal recovery problem (2.13), the latter reduces to a finite dimensional quadratic optimisation problem with linear constraints, as stated in the following corollary.

Corollary 4.4. *Let Assumption 3.10 hold for g . For all $n, m \in \mathbb{N}$, weights $\alpha_1, \dots, \alpha_{m+n} \in \mathbb{R}$ and collocation points $\omega_i, \dots, \omega_{m+n} \in \mathfrak{K}$, the optimal recovery problem (2.13) reduces to the following finite dimensional quadratic optimisation problem with linear constraints*

$$\min_{\alpha \in \mathbb{R}^{m+n}} \frac{1}{2} \alpha^\top \mathcal{K} \alpha \quad \text{subject to} \quad \tilde{\mathcal{K}} \alpha = \mathbf{b},$$

where $\tilde{\mathcal{K}} \in \mathbb{R}^{(m+n) \times (m+n)}$ and $\mathbf{b} \in \mathbb{R}^{m+n}$ are defined as follows

$$\tilde{\mathcal{K}} := \begin{pmatrix} \mathcal{L}_t \kappa(\omega^1, \omega^1) & \dots & \mathcal{L}_t \kappa(\omega^1, \omega^{m+n}) \\ \vdots & \ddots & \vdots \\ \mathcal{L}_t \kappa(\omega^m, \omega^1) & \dots & \mathcal{L}_t \kappa(\omega^m, \omega^{m+n}) \\ \kappa(\omega^{m+1}, \omega^1) & \dots & \kappa(\omega^{m+1}, \omega^{m+n}) \\ \vdots & \ddots & \vdots \\ \kappa(\omega^{m+n}, \omega^1) & \dots & \kappa(\omega^{m+n}, \omega^{m+n}) \end{pmatrix} \quad \text{and} \quad \mathbf{b} := \begin{pmatrix} g(\omega^1) \\ \vdots \\ g(\omega^m) \\ \phi(\omega^{m+1}) \\ \vdots \\ \phi(\omega^{m+n}) \end{pmatrix}.$$

This lifts the main computational obstacle and opens a clear path to solving the optimisation problem, thereby resolving Question (3) of 2.5.

Remark 4.5. The matrix $\tilde{\mathcal{K}}$ is well-defined by the differentiability of the product kernel.

4.2. Consistency. For any subset $\mathcal{N} \in \mathbb{N}_0^3$, let us define the $\mathcal{C}^\mathcal{N}(\mathfrak{K})$ norm as

$$\|h\|_\mathcal{N} := \sum_{(n_1, n_2, n_3) \in \mathcal{N}} \sup \left\{ |\partial_t^{n_1} \partial_x^{n_2} \partial_\gamma^{n_3} h(t, x, \gamma)| : (t, x, \gamma) \in \mathfrak{K}, \boldsymbol{\eta} \in \text{BV}_0^{\otimes n_3}, \|\boldsymbol{\eta}\|_1 \leq 1 \right\}.$$

Due to the fact that our estimates in Lemma A.1 do not hold for all $n \in \mathbb{N}$, we restrict to the specification needed in our examples

$$(4.6) \quad \mathcal{N} = \{(n_1, n_2, n_3) \in \mathbb{N}_0^3 : n_1 \leq 1, n_2 = n_3 = 0, \text{ or } n_1 = 0, n_2 + n_3 \leq 2\}.$$

We associate to it the range of derivatives $\mathcal{D} := \{\text{id}, \partial_t, \partial_x, \partial_x^2, \partial_\gamma^\eta \partial_x, \partial_\gamma^\eta, \partial_\gamma^{\eta\bar{\eta}}\}$ such that the norm also reads

$$\|h\|_\mathcal{N} := \sum_{\partial \in \mathcal{D}} \sup \left\{ |\partial h(\omega)| : \omega \in \mathfrak{K}, \eta, \bar{\eta} \in \text{BV}_0, \|\eta\|_1 \vee \|\bar{\eta}\|_{1-var} \leq 1 \right\}.$$

Our main result, inspired by [12, Theorem 1.2], follows. Its proof is given at the end of the section.

Theorem 4.6. *Let Assumption 3.10 hold for g . Let \mathcal{N} be as in (4.6) with $\eta, \bar{\eta} \in \text{BV}_0$ and assume either one of the following conditions holds:*

- (i) *The family $(u_{m,n})_{m,n \in \mathbb{N}}$ is relatively compact in $\mathcal{C}^\mathcal{N}(\mathfrak{K})$ and there exists a unique solution u^* to the PPDE (2.3) in $\mathcal{C}^\mathcal{N}(\mathfrak{K})$;*
- (ii) *There exists a unique solution u^* to the PPDE (2.3) in \mathcal{H} .*

If, moreover, as m and n tend to infinity, the collocation points $(\omega^i)_{i=1}^{m+n}$ form a countable dense family of \mathfrak{K} , then $(u_{m,n})_{m,n \in \mathbb{N}}$ converges towards u^ in the $\mathcal{C}^\mathcal{N}$ -topology as m, n tend to $+\infty$.*

The assumptions of the theorem all deserve separate discussion, studied in inverse order.

On the collocation points. It is natural to ask the collocation points to be dense in \mathfrak{K} , however this leaves a lot of flexibility for choosing them in an optimal way, especially in the space of paths. This issue is directly related to the rate of convergence of the method which lies beyond the scope of this paper. In practice though, one is only interested in covering the support of the evaluation points. During the course of our experiments, we found that sampling the collocation points randomly (as

Brownian motion trajectories) led to similar accuracy as sampling them from the same distribution as the evaluation points.

On the well-posedness of (2.3) Existence and uniqueness of (2.3) is discussed in several examples of interest in Section 2.3, and in the more general Volterra case in [10]. Uniqueness allows us to show that all the converging subsequences of $(u_{m,n})_{m,n}$ have the same limit, thus implying the sequence does converge. Unfortunately, known uniqueness results hold on a subset of $\mathcal{C}(\tilde{\Omega})$ whereas we require uniqueness in (a subset of) $\mathcal{C}(\mathfrak{K})$ where \mathfrak{K} is strictly included in $\tilde{\Omega}$.

On the assumption $u^* \in \mathcal{H}$.

- Proving convergence of $(u_{m,n})_{m,n}$ at the very least requires to find a norm in which this family is uniformly bounded. The reason for supposing $u^* \in \mathcal{H}$ is that it implies both the RKHS and $\mathcal{C}^{\mathcal{N}}$ norms of $(u_{m,n})_{m,n}$ are uniformly bounded, which in turn yields the relative compactness of this family in $\mathcal{C}^{\mathcal{N}}$, by Lemma 4.8. It is still slightly different from condition (i) because uniqueness holds in a different space.
- If $u^* \notin \mathcal{H}$ then the RKHS norm cannot be expected to be bounded and we lose our most promising tool to derive estimates for other norms. Note that we do not need equicontinuity if we look for compactness in \mathbb{R} , but we still require $u_{m,n}$ and their derivatives to be bounded uniformly in \mathbb{N}^2 .
- For u^* to be in \mathcal{H} would require at minima to be in \mathcal{C}^∞ . In this direction, Theorem 3.10 in [21] shows how to expand a path (semimartingale) functional with respect to terms of the signature. It is conceivable that such a representation also holds in the fractional setting for sufficiently smooth functionals, which could then be identified to an element of the signature RKHS. When the underlying process is a semimartingale, smoothness of the payoff function (and the coefficients) essentially ensures smoothness of the conditional expectation (i.e. the solution to the Kolmogorov equation). However when the direction of the pathwise derivative is singular (only square integrable) one can only prove twice differentiability of the solution (see [10, Proposition 2.23]). This remains a glass ceiling until one unveils how to exploit the regularisation properties of the expectation.

Remark 4.7. Several other works in the literature assume that $u^* \in \mathcal{H}$. [12] make in addition the classical assumption that \mathcal{H} is embedded in a Sobolev space, a condition we drop because we are able to prove sufficient regularity only with the help of the bounded RKHS norm. In finite dimensions, this condition allows to prove convergence rates, see [54, Proposition 11.30]. See also [5] who provide error bounds under additional assumptions.

The proof strategy for Theorem 4.6 using condition (ii) consists in showing that $(u_{m,n})$ is relatively compact and then extracting converging subsequences. Hence before presenting it we need to identify compact subsets of \mathcal{H} with respect to the $\mathcal{C}^{\mathcal{N}}$ -topology. The RKHS norm is linked to the regularity of the function, therefore a family of \mathcal{H} bounded in RKHS norm is equicontinuous. If this family is also bounded in $\mathcal{C}^{\mathcal{N}}$ norm we can conclude by Arzelà-Ascoli theorem that it is relatively compact.

Lemma 4.8. *Let Assumption 3.10 hold for g . Let $\mathcal{Y} \subset \mathcal{H}$ be bounded under both the RKHS and $\mathcal{C}^{\mathcal{N}}$ norms. Then \mathcal{Y} is relatively compact with respect to the $\mathcal{C}^{\mathcal{N}}$ topology.*

Proof. We will prove sequential compactness thanks to Arzelà-Ascoli theorem. Consider a sequence of functions $(h_n)_{n \in \mathbb{N}} \subset \mathcal{Y}$, bounded under the $\mathcal{C}^{\mathcal{N}}$ -topology and the RKHS norm by a constant C_h . We consider $\omega = (t, x, \gamma), \bar{\omega} = (t, x, \bar{\gamma}) \in \mathfrak{K}$, and we have by the reproducing property, Cauchy-Schwarz

inequality and (A.2) that

$$\begin{aligned}
|h_n(\omega) - h_n(\bar{\omega})| &= |\langle h_n, \Phi(\omega) - \Phi(\bar{\omega}) \rangle_{\mathcal{H}}| \\
&\leq \|h_n\|_{\mathcal{H}} \|\Phi(\omega) - \Phi(\bar{\omega})\|_{\mathcal{H}} \\
&\leq C_h \sqrt{k((t, x), (t, x))} \left(|\varphi(\gamma_t)| \|S(\gamma^g) - S(\bar{\gamma}^g)\|_{\overline{T(\mathcal{H}_g)}} + |\varphi(\gamma_t) - \varphi(\bar{\gamma}_t)| \|S(\bar{\gamma}^g)\|_{\overline{T(\mathcal{H}_g)}} \right) \\
&\leq C_h \sqrt{k_\infty} \|\gamma - \bar{\gamma}\|_{0;[t,T]} \left(\overline{DS_0}(|\gamma|_{1-var;[t,T]}, |\bar{\gamma}|_{1-var;[t,T]}) + \overline{S_0}(|\bar{\gamma}|_{1-var;[t,T]}) \right)
\end{aligned}$$

where $k_\infty := \sup_{(t,x,\gamma) \in \mathfrak{K}} k((t, x), (t, x)) < +\infty$. Since $\overline{S_0}, \overline{DS_0}$ is continuous and $|\gamma|_{1-var;[t,T]}, |\bar{\gamma}|_{1-var;[t,T]}$ are bounded by R , there exists a constant $C > 0$, independent of ω , such that $|h_n(\omega) - h_n(\bar{\omega})| \leq C \|\gamma - \bar{\gamma}\|_{0;[t,T]} \leq C \|\gamma - \bar{\gamma}\|_{p-var;[t,T]}$. Therefore, $\{h_n\}_{n \in \mathbb{N}}$ is equicontinuous on the compact \mathfrak{K} as it is equipped with the p -variation norm.

Following the approach of (4.1) and leveraging on (A.4) and (A.6), similar bounds hold for h_n 's derivatives with $\eta, \bar{\eta} \in \text{BV}(\mathbb{T}, V)$. For clarity, we assume that $\ell(\gamma_t, \bar{\gamma}_t) = 1$ (for the RBF kernel this corresponds to $\gamma_t = \bar{\gamma}_t$) which simplifies the estimates:

$$\begin{aligned}
|\partial_\gamma^\eta h_n(\omega) - \partial_\gamma^\eta h_n(\bar{\omega})| &\leq C_h \sqrt{k((t, x), (t, x))} \|\partial_\gamma^\eta S(\gamma^g|_{[t,T]}) - \partial_\gamma^\eta S(\bar{\gamma}^g|_{[t,T]})\|_{\overline{T(\mathcal{H}_g)}} \\
&\leq C_h \sqrt{k_\infty} \|\gamma - \bar{\gamma}\|_{0;[t,T]} \overline{DS_1}(|\gamma|_{1-var;[t,T]}, |\bar{\gamma}|_{1-var;[t,T]}, \|\eta\|_{1-var;[t,T]}); \\
|\partial_\gamma^{\eta\bar{\eta}} h_n(\omega) - \partial_\gamma^{\eta\bar{\eta}} h_n(\bar{\omega})| &\leq C_h \sqrt{k((t, x), (t, x))} \|\partial_\gamma^{\eta\bar{\eta}} S(\gamma^g|_{[t,T]}) - \partial_\gamma^{\eta\bar{\eta}} S(\bar{\gamma}^g|_{[t,T]})\|_{\overline{T(\mathcal{H}_g)}} \\
&\leq C_h \sqrt{k_\infty} \|\gamma - \bar{\gamma}\|_{0;[t,T]} \overline{DS_2}(|\gamma|_{1-var;[t,T]}, |\bar{\gamma}|_{1-var;[t,T]}, \|\eta\|_{1-var;[t,T]}, \|\bar{\eta}\|_{1-var;[t,T]}).
\end{aligned}$$

By the continuity of $\overline{DS_1}$ and $\overline{DS_2}$, we can also conclude that the families $(\partial_\gamma^\eta h_n)_{n \in \mathbb{N}}$ and $(\partial_\gamma^{\eta\bar{\eta}} h_n)_{n \in \mathbb{N}}$ are equicontinuous on the compact \mathfrak{K} with respect to the p -variation. The case where $\ell(\gamma_t, \bar{\gamma}_t) \neq 1$ would yield the same result by exploiting the boundedness and Lipschitz continuity of φ and its derivatives, and the estimates from Lemma A.1.

We only give details of the equicontinuity with respect to paths since the counterpart on $\mathbb{T} \times \mathbb{R}^d$ can be proved in a similar but easier fashion, using that the feature map associated to k and its two derivatives are bounded and Lipschitz continuous. Furthermore, equicontinuity of the derivatives with respect to time and space is also straightforward with the same arguments.

Coupled with the uniform $\mathcal{C}^\mathcal{N}$ bounds, Arzelà-Ascoli's theorem implies that $(\partial h_n)_{n \in \mathbb{N}}$ is relatively compact in $\mathcal{C}(\mathfrak{K})$ for all $\partial \in \mathcal{D}$. In particular, there exists a subsequence $(n_k)_{k \in \mathbb{N}}$ such that h_{n_k} converges as k goes to $+\infty$; then $(\partial_t h_{n_k})_{k \in \mathbb{N}}$ is a subset of a relatively compact space hence it is one itself and we can find a subsubsequence n_{k_l} such that $\partial_t h_{n_{k_l}}$ converges. We can go on for each further derivative until we have found a subsequence (n_k) for which ∂h_{n_k} converges in $\mathcal{C}(\mathfrak{K})$ converges for all $\partial \in \mathcal{D}$. This is precisely convergence in $\mathcal{C}^\mathcal{N}$. \square

Proof of Theorem 4.6. Assume that (ii) holds. Since u^* solves the PPDE (2.3) with the Fréchet derivative $\widehat{\partial}^\eta$ and that the Gateaux derivative $\widehat{\partial}^\eta$ is a weaker type, u^* also solves the weaker PPDE (2.12). In particular, u^* satisfies the constraints of the optimal recovery problem (4.3) at every point in \mathfrak{K} , and since $u_{m,n}$ is the minimiser, we must have $\|u_{m,n}\|_{\mathcal{H}} \leq \|u^*\|_{\mathcal{H}} < \infty$. Moreover, for all $\partial \in \mathcal{D}$ and $\eta, \bar{\eta} \in \text{BV}_0$ such that $\max(\|\eta\|_{1-var;[0,T]}, \|\bar{\eta}\|_{1-var;[0,T]}) \leq 1$, Cauchy-Schwarz inequality and the same calculations as in the proof of Proposition 4.1 yield

$$\sup_{\omega \in \mathfrak{K}} |\partial u_{m,n}(\omega)| = \sup_{\omega \in \mathfrak{K}} |\partial \langle u_{m,n}, \kappa(\omega, \cdot) \rangle_{\mathcal{H}}| \leq \sup_{\omega \in \mathfrak{K}} \|\partial \kappa(\omega, \cdot)\|_{\mathcal{H}} \|u_{m,n}\|_{\mathcal{H}}.$$

We take evaluation points over a compact and $\partial \kappa(\omega, \cdot)$ is continuous hence $\sup_{\omega \in \mathfrak{K}} \|\partial \kappa(\omega, \cdot)\|_{\mathcal{H}}^2 < \infty$. Thus there exists $C > 0$ such that for all such $\eta, \bar{\eta}$,

$$\sup_{\omega \in \mathfrak{K}} |\partial u_{m,n}(\omega)| \leq C \|u^*\|_{\mathcal{H}}, \quad \text{for all } \partial \in \mathcal{D}.$$

This proves that $(u_{m,n})_{m,n \in \mathbb{N}}$ is bounded under both the RKHS and the $\mathcal{C}^\mathcal{N}$ norms and therefore this family form a relatively compact space by Lemma 4.8, which means there exists a subsequence, also denoted $(u_{m,n})_{m,n \in \mathbb{N}}$ for conciseness, which converges in $\mathcal{C}^\mathcal{N}$ to a limit u_∞ , as $m, n \rightarrow \infty$. Since $(u_{m,n})_{m,n \in \mathbb{N}}$ is bounded under the RKHS norm, we have $u_\infty \in \mathcal{H}$.

Under condition (i), the same conclusion holds except that u_∞ is an element of the closure of \mathcal{H} with respect to the $\mathcal{C}^\mathcal{N}$ norm; the universality property of the RKHS entails that $\overline{\mathcal{H}} = \mathcal{C}^\mathcal{N}(\mathfrak{K})$. We now need to show that u_∞ solves the PPDE (2.3).

Let us define $v := \mathcal{L}_t u_\infty$ and $v_{m,n} := \mathcal{L}_t u_{m,n}$ for all $m, n \in \mathbb{N}$. For any $\omega \in \mathfrak{K}$, by the triangle inequality and because $v_{m,n}(\omega_i) = f(\omega_i)$ for any $1 \leq i \leq m$, we have

$$(4.7) \quad |v(\omega) - f(\omega)| \leq \min_{1 \leq i \leq m} \left\{ |v(\omega) - v(\omega_i)| + |v(\omega_i) - v_{m,n}(\omega_i)| + |f(\omega_i) - f(\omega)| \right\}.$$

Recall that v and f are uniformly continuous over the compact set \mathfrak{K} and, as m, n go to infinity, $(\omega_i)_{1 \leq i \leq m+n}$ form a dense family of \mathfrak{K} . Therefore, for all $\varepsilon > 0$ there exists $M, N \in \mathbb{N}$ such that, if $m \geq M$ and $n \geq N$, we have

$$\min_{1 \leq i \leq m} \left\{ |v(\omega) - v(\omega_i)| + |f(\omega_i) - f(\omega)| \right\} \leq \varepsilon.$$

In addition, $v_{m,n}$ converges uniformly to v because $u_{m,n}$ converges to u_∞ in $\mathcal{C}^\mathcal{N}$. Since $\varepsilon > 0$ was arbitrary, Equation (4.7) thus entails that $v(\omega) = f(\omega)$. Following a similar argument it can be shown that $u_\infty(T, x, \gamma) = \phi(T, x, \gamma)$ for any x, γ .

In conclusion, u_∞ is a solution of the PPDE (2.3). Under condition (i) (respectively (ii)), it belongs to $\mathcal{C}^\mathcal{N}$ (respectively \mathcal{H}) and is the unique classical solution in this space, implying that $u^* = u_\infty$. Since every convergent subsequence $(u_{m,n})_{m,n \in \mathbb{N}}$ converges to the same limit u^* , the whole sequence also converges to u^* in the $\mathcal{C}^\mathcal{N}$ topology. \square

5. NUMERICAL EXPERIMENTS

In this section, we present numerical experiments benchmarking the signature kernel PPDE solver presented in the previous section against either analytic solutions (if they are available) or classical Monte-Carlo solvers. We consider the two examples mentioned in the introduction, namely the path-dependent heat equation (2.7) and the rough Bergomi PPDE (2.9). For both experiments, we consider two types of errors between the predicted prices and true prices, namely the mean squared error (MSE), and the mean absolute error (MAE). The optimal kernel hyperparameters were determined by cross-validation. We will conclude the section with a discussion to compare our kernel approach with recent neural networks techniques for solving PPDEs. Our code is available at <https://github.com/crispitaigorico/sigppde>.

5.1. Fractional Brownian motion. In this first example, we consider a one dimensional fractional Brownian motion X with Hurst exponent $H = 0.1$. We recall that by [52, Theorem 4.1], under appropriate regularity assumptions on assumptions on the functions ϕ, f , the conditional expectation

$$\mathbb{E} \left[\phi(X_T) + \int_t^T f(s, X_s) ds \middle| \mathcal{F}_t \right] =: u(t, \Theta^t)$$

is realised as the solution of the path-dependent heat equation

$$\partial_t u(t, \omega) + \langle \partial_\omega^2 u(t, \omega), (K^t, K^t) \rangle + f(t, \omega) = 0, \quad u(T, \omega) = \phi(\omega_T).$$

In our experiments we choose $f \equiv 0$ and consider the following three instances of function ϕ , for which analytic expressions of the conditional expectation are available:

- (1) $f(x) = x$, $\mathbb{E}[f(X_T) | \mathcal{F}_t] = \Theta_T^t$;
- (2) $f(x) = e^{\nu x}$, $\mathbb{E}[f(X_T) | \mathcal{F}_t] = \exp \left(\nu \Theta_T^t + \frac{\nu^2 (T-t)^{2H}}{2} \right)$;

$$(3) \quad f(x) = (x - K)_+, \quad \mathbb{E}[f(X_T)|\mathcal{F}_t] = \frac{(T-t)^H}{\sqrt{2\pi}} \exp\left(-\frac{(K-\Theta_t^t)^2}{2(T-t)^{2H}}\right) - (K - \Theta_t^t)\phi\left(\frac{\Theta_t^t - K}{(T-t)^H}\right),$$

where $\nu, K \in \mathbb{R}$ and ϕ is the cumulative distribution function of the standard normal.

Because analytic prices are available, we limit ourselves to assess the performance of our kernel algorithm to recover these true prices. Collocation points for the kernel method were chosen uniformly on $[0, 1]$ for the time variable and sampled from the process Θ^t for the path variable.




FIGURE 1. Analytic prices against Signature Kernel prices with 400 collocation points.

As it can be observed from Figure 1, our kernel approach is capable to recovering with a high level of accuracy the prices for all considered payoff profiles.

5.2. Rough Bergomi. In this section, we make use of the following formulation of the rough Bergomi model introduced by [6] and already written down in the introduction

$$X_t := \int_0^t \sqrt{\psi_u(\widehat{W}_u^H)} dB_u - \frac{1}{2} \int_0^t \psi_u(\widehat{W}_u^H) du, \quad B_u := \rho W_u^1 + \sqrt{1 - \rho^2} W_u^2,$$

$$\psi_t(x) := \xi \exp\left\{\eta x - \frac{\eta^2}{2} t^{2H}\right\}, \quad \widehat{W}_t^H := \sqrt{2H} \int_0^t (t-u)^{H-\frac{1}{2}} dW_u^1,$$

for independent Brownian motions (W^1, W^2) . The *hybrid scheme* of [8] is used for efficient, $\mathcal{O}(\ell \log \ell)$, simulation of the Volterra process, Y_t^a , where ℓ is the length of the simulated paths. For the experiments we take the same parameters as in the repository², namely $T = 1$, $\xi = 0.055$, $\eta = 1.9$, $\rho = -0.9$. We set the ground truth to be the Monte Carlo prices obtained with 100000 sample paths.

$$(5.1) \quad \partial_t u + \frac{1}{2} \psi_t(\gamma) (\partial_{xx}^2 - \partial_x) u + \frac{1}{2} \langle \partial_{\gamma\gamma}^2 u, (K^t, K^t) \rangle + \rho \sqrt{\psi_t(\gamma)} \langle \partial_\gamma (\partial_x u), K^t \rangle = 0,$$

To evaluate the conditional expectation in the Monte Carlo benchmark we remark that, setting

$$I_s^t := \int_t^s K(s, r) dW_r, \quad \text{where } K(s, r) = \sqrt{2H} (s-r)^{H-\frac{1}{2}}$$

if $s > t$ and 0 otherwise, the conditional expectation reduces to the following unconditional expectation as we showed in the introduction

$$(5.2) \quad u(t, x, \omega) = \mathbb{E} \left[f \left(x + \int_t^T \sqrt{\psi_s(\omega_s + I_s^t)} dB_s - \frac{1}{2} \int_t^T \psi_s(\omega_s + I_s^t) ds \right) \right]$$

since Θ^t is an \mathcal{F}_t measurable process and I^t is independent from \mathcal{F}_t . Collocation points for the kernel method were sampled uniformly at random on $[0, 1]$ for the time variable, uniformly at random on $[x_{min}, x_{max}]$ for the price variables, where x_{min}, x_{max} were chosen using ground truth prices, and sampled from the process Θ^t for the path variable.

²https://github.com/ryanmccrickerd/rough_bergomi/tree/master.

$$H = 0.1 - \text{Strike} = 0.1$$

Model	Monte Carlo			SigPPDE		
Complexity	20 paths	100 paths	500 paths	20 c. pts	100 c. pts	500 c. pts
MSE	0.2070	0.0079	0.0080	0.2084	0.0016	0.0012
MAE	1.3193	0.2683	0.2300	0.6375	0.0714	0.0672

$$H = 0.1 - \text{Strike} = 1.0$$

Model	Monte Carlo			SigPPDE		
Complexity	20 paths	100 paths	500 paths	20 c. pts	100 c. pts	500 c. pts
MSE	0.0324	0.0013	0.0001	0.1576	0.0092	0.0008
MAE	0.0679	0.0557	0.0131	0.7998	0.2671	0.0938

$$H = 0.3 - \text{Strike} = 0.1$$

Model	Monte Carlo			SigPPDE		
Complexity	20 paths	100 paths	500 paths	20 c. pts	100 c. pts	500 c. pts
MSE	0.0198	0.0013	0.0005	0.1198	0.0026	0.0002
MAE	0.3330	0.0718	0.0624	0.6344	0.0973	0.0284

$$H = 0.3 - \text{Strike} = 1.0$$

Model	Monte Carlo			SigPPDE		
Complexity	20 paths	100 paths	500 paths	20 c. pts	100 c. pts	500 c. pts
MSE	0.0113	0.0015	0.0006	0.4201	0.0046	0.0014
MAE	0.2341	0.0945	0.0679	1.1832	0.1050	0.0905

We consider two values of $H = 0.1$ and $H = 0.3$ and two strike values $K = 0.1$ and $K = 1.0$. We take the same number of sample paths and collocation points for Monte Carlo and the signature kernel solver respectively, and consider three different increasing values, namely 20, 100 and 500. As it can be observed from the tables, the numerical performance in terms of MSE and MAE of the two algorithms are comparable across the different settings.

The time complexity of a Monte Carlo scheme scales roughly as $\mathcal{O}(h^{-2}Nd^2)$, where N is the number of sample paths, h is the discretisation step and d is the dimension of the process. This complexity can further be reduced to $\mathcal{O}((h \log(h))^{-1}Nd^2)$ using schemes such as the one proposed by [8]. The complexity for evaluating our solver (once it has been trained offline) scales as $\mathcal{O}(h^{-2}\hat{N}d^2)$ where \hat{N} is the number of collocation points. This cost can be further reduced $\mathcal{O}(h^{-1}\hat{N}d^2)$ if the operations are carried out on a GPU, see [44] for additional details. It is clear that a transparent comparison of complexities between our approach and Monte Carlo for pricing under rough volatility can only be achieved by obtaining error rates for the convergence of the approximations governing the values of N and \hat{N} respectively in the above complexities, which we intend to explore as future work.

Nevertheless, the PDE approach provides by design advantages over Monte Carlo methods. Indeed, finite-dimension greeks are prone to instability, let alone pathwise ones, while all derivatives of the PPDE solution appear from (the linear combination of) standard PDEs.

5.3. Comparison with neural PDE solvers. As anticipated, we will conclude this section with some remarks aimed at comparing our kernel approach to the numerous neural network techniques to solve PPDE which have been proposed in the literature in recent years. Neural PDE solvers such as DGM [49] and PINNs [41] essentially parameterise the solution of the PDE as a neural network which is then trained using the PDE and its boundary conditions as loss function. One drawback of these approaches in the path-dependent setting [43, 48, 31] is that they require a time-discretisation of the solution. The time grids over which solutions and derivatives are evaluated during training and testing need to agree, usually followed by an arbitrary interpolation. On the contrary, our kernel approach operates directly on continuous paths, and it is therefore mesh-free. Furthermore, derivatives can be accessed directly by differentiating kernels, with no need to prematurely discretise in time. We note that concomitant to this paper is [23], which uses a *neural rough differential equation* (Neural RDE) model [39, 46] to parameterise the solution of a PPDE and is therefore also mesh-free. Another drawback of neural PDE solvers is that their theoretical analysis is often limited to density/universal approximation results; showing the existence of a network of a requisite size achieving a certain error rate, without guarantees whether this network is computable in practice. In fact the optimisation, done with variants of stochastic gradient descent, is never convex so there is no guarantee to attain a global minimum. Our kernel approach boils down to a convex finite dimensional optimisation, guaranteed to attain the unique global minimum, and that is consistent with the original problem when the number of collocation points is sent to infinity.

6. OUTLOOK

The numerical solver presented in this paper competes with Monte Carlo methods in terms of the balance between complexity and accuracy. Our next objective will be to derive error estimates for the proposed method with respect to the number of collocation points in order to carry out a fair comparison. This step is essential for choosing the collocation points in an optimal way and better assess the efficiency of the algorithm. For this purpose we may follow the approach of [5] or decide to regularise our optimisation problem as in [15]. Further extensions include PPDEs arising from general Volterra processes, non-linear PPDEs, and path-dependent payoffs as the VIX example presented in Subsection 2.3.2. Finally, the inverse problem is particularly relevant for financial applications. We expect that the calibration of model parameters can be considerably sped up if the kernel is enhanced with this additional data.

APPENDIX A. PROOFS OF SIGNATURE AND SIGNATURE KERNEL ESTIMATES

A.1. Differentiability and continuity estimates for the signature. For the sake of uncluttering notations, in this section we fix $t \in \mathbb{T}$ and we will write $|\gamma|_1$ instead of $|\gamma|_{1-var;[t,T]}$ and $\|\gamma\|_0$ instead of $\|\gamma\|_{0;[t,T]}$ for every $\gamma \in BV_t$. We present estimates for the signature and its derivatives, in passing obtaining CDEs satisfied by the latter, and showing continuity with respect to the C^0 -norm as claimed in Proposition 3.14. The following lemma is inspired from Theorems 3.15 and 4.4 from [25].

Lemma A.1. *Let Assumption 3.10 hold for g . For all $\gamma, \bar{\gamma}, \eta, \bar{\eta} \in BV_t$, the signature is two times differentiable. Moreover, there exist continuous functions $\overline{S}_k : \mathbb{R}_+^k \rightarrow \mathbb{R}_+$ and $\overline{DS}_k : \mathbb{R}_+^k \rightarrow \mathbb{R}_+$,*

for $k = 1, 2, 3$, such that the following estimates hold

$$\begin{aligned}
(\text{A.1}) \quad & \sup_{t' \in [t, T]} \|S(\gamma^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \leq \overline{S_0}(|\gamma|_1); \\
(\text{A.2}) \quad & \sup_{t' \in [t, T]} \|S(\gamma^g)_{t'} - S(\bar{\gamma}^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \leq \|\gamma - \bar{\gamma}\|_0 \overline{DS_0}(|\gamma|_1, |\bar{\gamma}|_1); \\
(\text{A.3}) \quad & \sup_{t' \in [t, T]} \|\partial_\gamma^\eta S(\gamma^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \leq \overline{S_1}(|\gamma|_1, \|\eta\|_1); \\
(\text{A.4}) \quad & \sup_{t' \in [t, T]} \|\partial_\gamma^\eta S(\gamma^g)_{t'} - \partial_\gamma^\eta S(\bar{\gamma}^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \leq \|\gamma - \bar{\gamma}\|_0 \overline{DS_1}(|\gamma|_1, |\bar{\gamma}|_1, \|\eta\|_1); \\
(\text{A.5}) \quad & \sup_{t' \in [t, T]} \|\partial_\gamma^{\eta\bar{\eta}} S(\gamma^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \leq \overline{S_2}(|\gamma|_1, \|\eta\|_1, \|\bar{\eta}\|_1); \\
(\text{A.6}) \quad & \sup_{t' \in [t, T]} \|\partial_\gamma^{\eta\bar{\eta}} S(\gamma^g)_{t'} - \partial_\gamma^{\eta\bar{\eta}} S(\bar{\gamma}^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \leq \|\gamma - \bar{\gamma}\|_0 \overline{DS_2}(|\gamma|_1, |\bar{\gamma}|_1, \|\eta\|_1, \|\bar{\eta}\|_1).
\end{aligned}$$

Proof. Let $0 \leq t \leq t' \leq T$ and $\gamma, \bar{\gamma}, \eta, \bar{\eta} \in \text{BV}_t$.

1) The signature satisfies

$$S(\gamma^g)_{t'} = \mathbf{1} + \int_t^{t'} S(\gamma)_u \cdot d\gamma_u^g = \mathbf{1} + \int_t^{t'} S(\gamma)_u \cdot \nabla g(\gamma_u, \cdot) d\gamma_u,$$

hence we have the bound

$$\begin{aligned}
(\text{A.7}) \quad & \|S(\gamma^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \leq \mathbf{1} + \int_t^{t'} \|S(\gamma)_u\|_{\overline{T(\mathcal{H}_g)}} \cdot \|\nabla g(\gamma_u, \cdot) d\gamma_u\|_{\overline{T(\mathcal{H}_g)}} \\
& \leq \mathbf{1} + \int_t^{t'} \|S(\gamma)_u\|_{\overline{T(\mathcal{H}_g)}} \|\nabla g(\gamma_u, \cdot)\|_{\mathcal{H}_g} |d\gamma_u|,
\end{aligned}$$

where $\|\nabla g(\gamma_u, \cdot)\|_{\mathcal{H}_g} = \sum_{i,j=1}^e \partial_{x^i y^j}^2 g(\gamma_u, \gamma_u) \leq \overline{g_{11}}$ and $|d\gamma_u| = |\dot{\gamma}_u| du$. By Grönwall's lemma [25, Lemma 3.2] this yields

$$\sup_{t' \in [t, T]} \|S(\gamma^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \leq e^{\overline{g_{11}}|\gamma|_1} =: \overline{S_0}(\gamma).$$

2) Furthermore, [25, Proposition 2.9] gives the 1-variation bound

$$\left\| \|S(\gamma^g)_{\cdot}\|_{\overline{T(\mathcal{H}_g)}} \right\|_1 \leq \overline{S_0}(\gamma) \overline{g_{11}} |\gamma|_1.$$

For any $\bar{\gamma} \in \text{BV}_t$, imitating the proof of [25, Theorem 3.15] yields

$$\begin{aligned}
\|S(\gamma^g)_{t'} - S(\bar{\gamma}^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} & \leq \int_t^{t'} \|S(\gamma^g)_u - S(\bar{\gamma}^g)_u\|_{\overline{T(\mathcal{H}_g)}} |d\gamma_u^g| + \left\| \int_t^{t'} S(\bar{\gamma}^g) d(\gamma^g - \bar{\gamma}^g)_u \right\|_{\overline{T(\mathcal{H}_g)}} \\
& \leq \int_t^{t'} \|S(\gamma^g)_u - S(\bar{\gamma}^g)_u\|_{\overline{T(\mathcal{H}_g)}} \overline{g_{11}} |d\gamma_u| + \left\| \gamma^g - \bar{\gamma}^g \right\|_{\mathcal{H}_g} \left(1 + \left\| \|S(\bar{\gamma}^g)_{\cdot}\|_{\overline{T(\mathcal{H}_g)}} \right\|_1 \right)
\end{aligned}$$

Thus, applying the inequality $\left\| \gamma^g - \bar{\gamma}^g \right\|_{\mathcal{H}_g} \leq \overline{g_{11}} \|\gamma - \bar{\gamma}\|_0$ and Grönwall's lemma we obtain

$$\sup_{t' \in [t, T]} \|S(\gamma^g)_{t'} - S(\bar{\gamma}^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \leq \|\gamma - \bar{\gamma}\|_0 \left(1 + \overline{S_0}(\bar{\gamma}) \overline{g_{11}} |\bar{\gamma}|_1 \right) e^{\overline{g_{11}}|\gamma|_1} =: \|\gamma - \bar{\gamma}\|_0 \overline{DS_0}(|\gamma|_1, |\bar{\gamma}|_1),$$

where $\overline{DS_0} : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a continuous function.

3) We start with the following observation. For all $u \in [t, T]$ we have

$$\begin{aligned}
\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} d(\gamma + \varepsilon \eta)_u^g - \gamma_u^g &= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \sum_{i=1}^e \left\{ \partial_i g(\gamma_u + \varepsilon \eta_u, \cdot) d(\gamma_u^i + \varepsilon \eta_u^i) - \partial_i g(\gamma_u, \cdot) d\gamma_u^i \right\} \\
&= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \sum_{i=1}^e \left\{ \partial_i g(\gamma_u + \varepsilon \eta_u, \cdot) - \partial_i g(\gamma_u, \cdot) \right\} d\gamma_u^i + \lim_{\varepsilon \rightarrow 0} \sum_{i=1}^e \partial_i g(\gamma_u + \varepsilon \eta_u, \cdot) d\eta_u^i \\
&= \sum_{i=1}^e \left\{ \partial_\gamma^\eta \partial_{x^i} g(\gamma_u, \cdot) d\gamma_u^i + \partial_{x^i} g(\gamma_u, \cdot) d\eta_u^i \right\} \\
&= d \left\{ \sum_{i=1}^e \partial_{x^i} g(\gamma_u, \cdot) \eta_u^i \right\} \\
\text{(A.8)} \quad &= d \left\{ \partial_\gamma^\eta G(\gamma)_u \right\}
\end{aligned}$$

Further, we note that, for all $\varepsilon > 0$,

$$\frac{1}{\varepsilon} \sup_{t' \in [t, T]} \|S(\gamma^g)_{t'} - S((\gamma + \varepsilon \eta)^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \leq \|\eta\|_0 \overline{DS_0}(|\gamma|_1, |\gamma|_1 + |\eta|_1).$$

Hence, by dominated convergence and (A.8), we get

$$\begin{aligned}
\partial_\gamma^\eta S(\gamma^g)_{t'} &= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} (S((\gamma + \varepsilon \eta)^g)_{t'} - S(\gamma^g)_{t'}) \\
&= \int_t^{t'} \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} (S((\gamma + \varepsilon \eta)^g)_u - S(\gamma^g)_u) \cdot d\gamma_u^g + \int_t^{t'} \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} S(\gamma^g)_u \cdot d((\gamma + \varepsilon \eta)^g)_u - \gamma_u^g \\
\text{(A.9)} \quad &= \int_t^{t'} \partial_\gamma^\eta S(\gamma^g)_u \cdot d\gamma_u^g + \int_t^{t'} S(\gamma^g)_u \cdot d \left\{ \partial_\gamma^\eta G(\gamma)_u \right\}.
\end{aligned}$$

Note that $\|\partial_{x^{ij}}^2 g(\gamma_u, \cdot)\|_{\mathcal{H}_g} = \partial_{x^{ij} y^{ij}}^4 g(\gamma_u, \gamma_u)$, and hence

$$\begin{aligned}
\int_t^{t'} \|d \left\{ \partial_\gamma^\eta G(\gamma)_u \right\}\|_{\mathcal{H}_g} &\leq \int_t^{t'} \sum_{i,j=1}^e \|\partial_{x^{ij}} g(\gamma_u, \cdot)\|_{\mathcal{H}_g} |\eta_u| |d\gamma_u| + \int_t^{t'} \sum_{i=1}^e \|\partial_{x^i} g(\gamma_u, \cdot)\|_{\mathcal{H}_g} |d\eta_u| \\
\text{(A.10)} \quad &\leq \overline{g_{22}} \|\eta\|_0 |\gamma|_1 + \overline{g_{11}} |\eta|_1 \leq \overline{G_1}(|\gamma|_1, \|\eta\|_1),
\end{aligned}$$

where $\overline{G_1} : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ is a continuous function. Applying Grönwall's lemma yields the estimate

$$\|\partial_\gamma^\eta S(\gamma^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \leq \overline{S_0}(|\gamma|_1) \left(\overline{g_{22}} \|\eta\|_0 |\gamma|_1 + \overline{g_{11}} |\eta|_1 \right) e^{\overline{g_{11}} |\gamma|_1} =: \overline{S_1}(|\gamma|_1, \|\eta\|_1).$$

4) Noting that, for all $t \leq s' \leq t' \leq T$,

$$\left\| \int_t^{t'} S(\gamma^g)_u \cdot d \left\{ \partial_\gamma^\eta G(\gamma)_u \right\} \right\|_{\overline{T(\mathcal{H}_g)}} \leq \overline{S_0}(\gamma) \left\| \int_t^{t'} d \left\{ \partial_\gamma^\eta G(\gamma)_u \right\} \right\|_{\mathcal{H}_g} \leq \overline{S_0}(\gamma) \overline{G_1}(|\gamma|_1, \|\eta\|_1),$$

Proposition 2.9 of [25] thus gives

$$\left\| \partial_\gamma^\eta S(\gamma^g)_{t'} \right\|_{\overline{T(\mathcal{H}_g)}} \leq \overline{S_1}(|\gamma|_1, \|\eta\|_1) \overline{g_{11}} |\gamma|_1 + \overline{S_0}(\gamma) \overline{G_1}(|\gamma|_1, \|\eta\|_1).$$

Then, from the integration by parts formula we observe that

$$\int_t^{t'} |\eta_u| d(|\gamma_u - \bar{\gamma}_u|) = \int_t^{t'} |\gamma_u - \bar{\gamma}_u| d|\eta_u| + |\eta_t| |\gamma_t - \bar{\gamma}_t| \leq \|\gamma - \bar{\gamma}\|_0 |\eta|_1 + \|\eta\|_0 \|\gamma - \bar{\gamma}\|_0.$$

This allows us to study the following regularity:

$$\begin{aligned}
& \int_t^{t'} \left\| d \left\{ \partial_\gamma^\eta (G(\gamma)_u - G(\bar{\gamma})_u) \right\} \right\|_{\mathcal{H}_g} \leq \int_t^{t'} \sum_{i,j=1}^e \left\| \partial_{x^{ij}} (g(\gamma_u, \cdot) - g(\bar{\gamma}_u, \cdot)) \right\|_{\mathcal{H}_g} |\eta_u| \, d\gamma_u \\
& + \int_t^{t'} \sum_{i,j=1}^e \left\| \partial_{x^{ij}} g(\bar{\gamma}_u, \cdot) \right\|_{\mathcal{H}_g} |\eta_u| \, d(\gamma_u - \bar{\gamma}_u) + \int_t^{t'} \sum_{i=1}^e \left\| \partial_{x^i} (g(\gamma_u, \cdot) - g(\bar{\gamma}_u, \cdot)) \right\|_{\mathcal{H}_g} \, d\eta_u \\
& \leq \overline{g_{33}} \|\gamma - \bar{\gamma}\|_0 \|\eta\|_0 |\gamma|_1 + \overline{g_{22}} \int_t^{t'} |\eta_u| \, d(|\gamma_u - \bar{\gamma}_u|) + \overline{g_{22}} \|\gamma - \bar{\gamma}\|_0 \|\eta\|_1 \\
& \leq \|\gamma - \bar{\gamma}\|_0 \overline{DG_1}(|\gamma|_1, |\bar{\gamma}|_1, \|\eta\|_1),
\end{aligned}$$

where $\overline{DG_1} : \mathbb{R}_+^3 \rightarrow \mathbb{R}_+$ is a continuous function. We deduce that the difference of the source terms in (A.9) has the following bound

$$\begin{aligned}
& \left\| \int_t^{t'} S(\gamma^g)_u \cdot d \left\{ \partial_\gamma^\eta G(\gamma)_u \right\} - \int_t^{t'} S(\bar{\gamma}^g)_u \cdot d \left\{ \partial_\gamma^\eta G(\bar{\gamma})_u \right\} \right\|_{\overline{T(\mathcal{H}_g)}} \\
& \leq \sup_{t' \in [t, T]} \|S(\gamma^g)_{t'} - S(\bar{\gamma}^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \overline{G_1}(\gamma, \eta) + \sup_{t' \in [t, T]} \|S(\gamma^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \int_t^{t'} \left\| d \left\{ \partial_\gamma^\eta (G(\gamma)_u - G(\bar{\gamma})_u) \right\} \right\|_{\mathcal{H}_g} \\
& \leq \|\gamma - \bar{\gamma}\|_0 \left(\overline{DS_0}(|\gamma|_1, |\bar{\gamma}|_1) \overline{G_1}(|\gamma|_1, \|\eta\|_1) + \overline{S_0}(\gamma) \overline{DG_1}(|\gamma|_1, |\bar{\gamma}|_1, \|\eta\|_1) \right).
\end{aligned}$$

In virtue of (A.9) and (A.11), the same arguments as in the proof of [25, Theorem 3.15] entails there exists a continuous function $\overline{DS_1} : \mathbb{R}_+^3 \rightarrow \mathbb{R}_+$ such that

$$\begin{aligned}
& \sup_{t' \in [t, T]} \left\| \partial_\gamma^\eta S(\gamma^g)_{t'} - \partial_\gamma^\eta S(\bar{\gamma}^g)_{t'} \right\|_{\overline{T(\mathcal{H}_g)}} \\
& \leq e^{2\overline{g_{11}}(|\gamma|_1 \vee |\bar{\gamma}|_1)} \left\| \int_t^{t'} S(\gamma^g)_u \cdot d \left\{ \partial_\gamma^\eta G(\gamma)_u \right\} - \int_t^{t'} S(\bar{\gamma}^g)_u \cdot d \left\{ \partial_\gamma^\eta G(\bar{\gamma})_u \right\} \right\|_{\overline{T(\mathcal{H}_g)}} \\
& \leq \|\gamma - \bar{\gamma}\|_0 \overline{DS_1}(|\gamma|_1, |\bar{\gamma}|_1, \|\eta\|_1).
\end{aligned}$$

5) The latest estimate allows to use dominated convergence one more time and compute the second derivative in the direction of $\eta, \bar{\eta}$:

$$\partial_\gamma^{\eta \bar{\eta}} S(\gamma^g)_{t'} = \int_t^{t'} \partial_\gamma^{\eta \bar{\eta}} S(\gamma^g)_u \, d\gamma_u^g + \text{Source}(\gamma, \eta, \bar{\eta}),$$

where

$$\text{Source}(\gamma, \eta, \bar{\eta}) := \int_t^{t'} \partial_\gamma^\eta S(\gamma^g)_u \cdot d \left\{ \partial_{\bar{\gamma}}^{\bar{\eta}} G(\gamma)_u \right\} + \int_t^{t'} \partial_{\bar{\gamma}}^{\bar{\eta}} S(\gamma^g)_u \cdot d \left\{ \partial_\gamma^\eta G(\gamma)_u \right\} + \int_t^{t'} S(\gamma^g)_u \cdot d \left\{ \partial_\gamma^{\eta \bar{\eta}} G(\gamma)_u \right\}.$$

Using the same arguments as in (A.8) one deduces that

$$\begin{aligned}
& \int_t^{t'} \left\| d \left\{ \partial_\gamma^{\eta \bar{\eta}} G(\gamma)_u \right\} \right\|_{\mathcal{H}_g} \leq \int_t^{t'} \sum_{i,j,k=1}^e \left\| \partial_{x^{ijk}}^3 g(\gamma_u, \cdot) \right\|_{\mathcal{H}_g} |\bar{\eta}_u| |\eta_u| \, d\gamma_u \\
& + \int_t^{t'} \sum_{i,j=1}^e \left\| \partial_{x^{ij}}^2 g(\gamma_u, \cdot) \right\|_{\mathcal{H}_g} \left(|\eta_u| \, d\bar{\eta}_u + |\bar{\eta}_u| \, d\eta_u \right) \\
& \leq \overline{g_{33}} \|\eta\|_0 \|\bar{\eta}\|_0 |\gamma|_1 + \overline{g_{22}} \left(\|\eta\|_0 \|\bar{\eta}\|_1 + \|\bar{\eta}\|_0 \|\eta\|_1 \right) \leq \overline{G_2}(|\gamma|_1, \|\eta\|_1, \|\bar{\eta}\|_1),
\end{aligned}$$

where $\overline{G}_2 : \mathbb{R}_+^3 \rightarrow \mathbb{R}_+$ is a continuous function. Grönwall's inequality once again yields the bound

$$\sup_{t' \in [t, T]} \|\partial_{\gamma}^{\eta\bar{\eta}} S(\gamma^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \leq \left(\overline{S}_1(\gamma, \eta) \overline{G}_1(\gamma, \bar{\eta}) + \overline{S}_1(\gamma, \bar{\eta}) \overline{G}_1(\gamma, \eta) + \overline{S}_0(\gamma) \overline{G}_2(\gamma, \eta, \bar{\eta}) \right) e^{\overline{S}_1 |\gamma|_1}.$$

6) The same steps as for the first derivative yield the existence of a continuous function $\overline{DG}_2 : \mathbb{R}^4 \rightarrow \mathbb{R}$ such that

$$\int_t^{t'} \|\mathrm{d} \{ \partial_{\gamma}^{\eta\bar{\eta}} (G(\gamma)_u - G(\bar{\gamma})_u) \}\|_{\mathcal{H}_g} \leq \|\gamma - \bar{\gamma}\|_0 \overline{DG}_2(|\gamma|_1, |\bar{\gamma}|_1, \|\eta\|_1, \|\bar{\eta}\|_1).$$

Hence the difference of the source term for two paths $\gamma, \bar{\gamma}$ can be bounded as follows

$$\begin{aligned} & \|\text{Source}(\gamma, \eta, \bar{\eta}) - \text{Source}(\bar{\gamma}, \eta, \bar{\eta})\|_{\overline{T(\mathcal{H}_g)}} \\ & \leq \|\gamma - \bar{\gamma}\|_0 \left(\overline{S}_1(|\gamma|_1, \|\eta\|_1) \overline{DG}_1(|\gamma|_1, |\bar{\gamma}|_1, \|\eta\|_1) + \overline{G}_1(|\gamma|_1, |\bar{\gamma}|_1, \|\eta\|_1) \overline{DS}_1(|\gamma|_1, |\bar{\gamma}|_1, \|\eta\|_1) \right. \\ & \quad + \overline{S}_1(|\gamma|_1, \|\bar{\eta}\|_1) \overline{DG}_1(|\gamma|_1, |\bar{\gamma}|_1, \|\bar{\eta}\|_1) + \overline{G}_1(|\gamma|_1, |\bar{\gamma}|_1, \|\bar{\eta}\|_1) \overline{DS}_1(|\gamma|_1, |\bar{\gamma}|_1, \|\bar{\eta}\|_1) \\ & \quad \left. + \overline{S}_0(|\gamma|_1) \overline{G}_2(|\gamma|_1, |\bar{\gamma}|_1, \|\eta\|_1, \|\bar{\eta}\|_1) + \overline{DS}_0(|\gamma|_1, |\bar{\gamma}|_1) \overline{G}_2(|\gamma|_1, |\bar{\gamma}|_1, \|\eta\|_1, \|\bar{\eta}\|_1) \right). \end{aligned}$$

Finally, by [25, Theorem 3.15] again there is a continuous function $\overline{DS}_2 : \mathbb{R}^4 \rightarrow \mathbb{R}$ such that

$$\sup_{t' \in [t, T]} \|\partial_{\gamma}^{\eta\bar{\eta}} S(\gamma^g)_{t'} - \partial_{\gamma}^{\eta\bar{\eta}} S(\bar{\gamma}^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \leq \|\gamma - \bar{\gamma}\|_0 \overline{DS}_2(|\gamma|_1, |\bar{\gamma}|_1, \|\eta\|_1, \|\bar{\eta}\|_1)$$

thus concluding the proof. \square

A.2. Differentiability and continuity estimates for the signature kernel. We start by stating estimates for the signature kernel and its derivatives, which follow from Lemma A.1. The PDE formulations in Proposition 3.11 are immediate corollaries. In terms of notations, for two paths $\gamma \in \text{BV}_t$ and $\tau \in \text{BV}_s$, we will denote $|\gamma|_1 = |\gamma|_{1-\text{var};[t, T]}$ and $|\tau|_1 = |\tau|_{1-\text{var};[s, T]}$.

Lemma A.2. *Let Assumption and 3.10 hold for g . Let $s, t \in \mathbb{T}$, the paths $\gamma, \eta, \bar{\eta} \in \text{BV}_t$ and $\tau \in \text{BV}_s$. Then the following estimates hold:*

$$\begin{aligned} & \sup_{s' \in [s, T], t' \in [t, T]} |\kappa^g(\gamma, \tau)_{s', t'}| \leq \overline{S}_0(|\gamma|_1) \overline{S}_0(|\tau|_1); \\ & \sup_{s' \in [s, T], t' \in [t, T]} |\partial_{\gamma}^{\eta} \kappa^g(\gamma, \tau)_{s', t'}| \leq \overline{S}_1(|\gamma|_1, \|\eta\|_1) \overline{S}_0(|\tau|_1); \\ & \sup_{s' \in [s, T], t' \in [t, T]} |\partial_{\gamma}^{\eta\bar{\eta}} \kappa^g(\gamma, \tau)_{s', t'}| \leq \overline{S}_2(|\gamma|_1, \|\eta\|_1, \|\bar{\eta}\|_1) \overline{S}_0(|\tau|_1). \end{aligned}$$

Proof. Let $0 \leq s \leq s' \leq T$, $0 \leq t \leq t' \leq T$ and $\gamma, \eta, \bar{\eta} \in \text{BV}_t, \tau \in \text{BV}_s$. By the definition of the signature kernel, Cauchy-Schwarz inequality and (A.1) we have

$$|\kappa^g(\gamma, \tau)_{s', t'}| = |\langle S(\gamma^g)_{t'}, S(\tau^g)_{s'} \rangle| \leq \|S(\gamma^g)_{t'}\|_{\overline{T(\mathcal{H}_g)}} \|S(\tau^g)_{s'}\|_{\overline{T(\mathcal{H}_g)}} \leq \overline{S}_0(|\gamma|_1) \overline{S}_0(|\tau|_1).$$

Then for any $\varepsilon > 0$, Cauchy-Schwarz inequality and (A.2) yield

$$|\kappa^g(\gamma + \varepsilon\eta, \tau)_{s', t'} - \kappa^g(\gamma, \tau)_{s', t'}| \leq \|\eta\|_0 \overline{DS}_0(|\gamma|_1, |\gamma|_1 + |\eta|_1) \overline{S}_0(|\tau|_1).$$

Therefore, dominated convergence, Cauchy-Schwarz and (A.3) entail

$$|\partial_{\gamma}^{\eta} \kappa^g(\gamma, \tau)_{s', t'}| = |\langle \partial_{\gamma}^{\eta} S(\gamma^g)_{t'}, S(\tau^g)_{s'} \rangle| \leq \overline{S}_1(|\gamma|_1, \|\eta\|_1) \overline{S}_0(|\tau|_1).$$

The same ideas coupled with the bounds (A.5) and (A.6) lead to the last estimate. \square

Proof of Proposition 3.11. Let $0 \leq s \leq s' \leq T$ and $0 \leq t \leq t' \leq T$ and $\gamma, \eta, \bar{\eta} \in \text{BV}_t, \tau \in \text{BV}_s$.

1) In virtue of (A.2), dominated convergence allows to push the limit inside the integrals:

$$\begin{aligned}
\partial_\gamma^\eta \kappa^g(\gamma, \tau)_{s', t'} &= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} (\kappa^g(\gamma + \varepsilon \eta, \tau)_{s', t'} - \kappa^g(\gamma, \tau)_{s', t'}) \\
&= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \int_s^{s'} \int_t^{t'} \left(\kappa^g(\gamma + \varepsilon \eta, \tau)_{u,v} \langle d(\gamma + \varepsilon \eta)_u^g, d\tau_v^g \rangle_{\mathcal{H}_g} - \kappa^g(\gamma, \tau)_{u,v} \langle d\gamma_u^g, d\tau_v^g \rangle_{\mathcal{H}_g} \right) \\
&= \int_s^{s'} \int_t^{t'} \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \left(\kappa^g(\gamma + \varepsilon \eta, \tau)_{u,v} - \kappa^g(\gamma, \tau)_{u,v} \right) \langle d(\gamma + \varepsilon \eta)_u^g, d\tau_v^g \rangle_{\mathcal{H}_g} \\
&\quad + \kappa^g(\gamma, \tau)_{u,v} \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \langle d(\gamma + \varepsilon \eta)_u^g - d\gamma_u^g, d\tau_v^g \rangle_{\mathcal{H}_g} \\
&= \int_s^{s'} \int_t^{t'} \left(\partial_\gamma^\eta \kappa^g(\gamma, \tau)_{u,v} \langle d\gamma_u^g, d\tau_v^g \rangle_{\mathcal{H}_g} + \kappa^g(\gamma, \tau)_{u,v} \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \langle d(\gamma + \varepsilon \eta)_u^g - d\gamma_u^g, d\tau_v^g \rangle_{\mathcal{H}_g} \right).
\end{aligned}$$

For all $u \in [t, T]$ and $v \in [s, T]$ we have

$$\begin{aligned}
&\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \langle d(\gamma + \varepsilon \eta)_u^g - \gamma_u^g, d\tau_v^g \rangle_{\mathcal{H}_g} \\
&= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \sum_{i,k=1}^e \left\langle \partial_i g(\gamma_u + \varepsilon \eta_u, \cdot) d(\gamma_u^i + \varepsilon \eta_u^i) - \partial_i g(\gamma_u, \cdot) d\gamma_u^i, \partial_k g(\tau_v, \cdot) d\tau_v^k \right\rangle \\
&= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \sum_{i,k=1}^e \left\langle \partial_i g(\gamma_u + \varepsilon \eta_u, \cdot) - \partial_i g(\gamma_u, \cdot), \partial_k g(\tau_v, \cdot) \right\rangle d\gamma_u^i d\tau_v^k + \lim_{\varepsilon \rightarrow 0} \sum_{i,k=1}^e \left\langle \partial_i g(\gamma_u + \varepsilon \eta_u, \cdot), \partial_k g(\tau_v, \cdot) \right\rangle d\eta_u^i d\tau_v^k \\
&= \sum_{i,k=1}^e \lim_{\varepsilon \rightarrow 0} \left\{ \frac{1}{\varepsilon} \partial_{x^i} \partial_{y^k} \left(g(\gamma_u + \varepsilon \eta_u, \tau_v) - g(\gamma_u, \tau_v) \right) d\gamma_u^i d\tau_v^k + \partial_{x^i} \partial_{y^k} g(\gamma_u + \varepsilon \eta_u, \tau_v) d\eta_u^i d\tau_v^k \right\} \\
&= \sum_{i,k=1}^e \left\{ \partial_\gamma^\eta \partial_{x^i y^k}^2 g(\gamma_u, \tau_v) d\gamma_u^i + \partial_{x^i y^k}^2 g(\gamma_u, \tau_v) d\eta_u^i \right\} d\tau_v^k \\
&\quad \text{(A.14)} \\
&= \sum_{i,k=1}^e \left\{ \sum_{j=1}^e \partial_{x^{ij} y^k}^3 g(\gamma_u, \tau_v) \eta_u^j d\gamma_u^i + \partial_{x^i y^k}^2 g(\gamma_u, \tau_v) d\eta_u^i \right\} d\tau_v^k
\end{aligned}$$

where we used the reproducing property of derivatives of g (3.13) and we can swap partial derivatives because $g(\cdot, x) \in C^2(V, \mathbb{R})$ for all $x \in V$. Unrolling in the other direction one can also obtain

$$\begin{aligned}
\sum_{i,k=1}^e \left\{ \sum_{j=1}^e \partial_{x^{ij} y^k}^3 g(\gamma_u, \tau_v) \eta_u^j d\gamma_u^i + \partial_{x^i y^k}^2 g(\gamma_u, \tau_v) d\eta_u^i \right\} d\tau_v^k &= \left\langle d \left\{ \sum_{i=1}^e \partial_{x^i} g(\gamma_u, \cdot) \eta_u^i \right\}, d\tau_v^g \right\rangle \\
&= \left\langle d \left\{ \partial_\gamma^\eta G(\gamma)_u \right\}, d\tau_v^g \right\rangle.
\end{aligned}
\tag{A.15}$$

The expressions (A.15) and (A.14) yield (3.16) and (3.18) respectively.

2) We take the terms of (3.18) one by one to compute

$$\begin{aligned}
\partial_{\gamma}^{\eta\bar{\eta}} \kappa^g(\gamma, \tau)_{s', t'} &:= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \left(\partial_{\gamma}^{\eta} \kappa^g(\gamma + \varepsilon \bar{\eta}, \tau)_{s', t'} - \partial_{\gamma}^{\eta} \kappa^g(\gamma, \tau)_{s', t'} \right) \\
&= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \left\{ \int_s^{s'} \int_t^{t'} \partial_{\gamma}^{\eta} \kappa^g(\gamma + \varepsilon \bar{\eta}, \tau)_{u,v} \sum_{i,k=1}^e \partial_{x^i y^k}^2 g(\gamma_u + \varepsilon \bar{\eta}_u, \tau_v) d(\gamma_u^i + \varepsilon \bar{\eta}_u^i) d\tau_v^k \right. \\
&\quad \left. - \int_s^{s'} \int_t^{t'} \partial_{\gamma}^{\eta} \kappa^g(\gamma, \tau)_{u,v} \sum_{i,k=1}^e \partial_{x^i y^k}^2 g(\gamma_u, \tau_v) d\gamma_u^i d\tau_v^k \right\} \\
&+ \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \left\{ \int_s^{s'} \int_t^{t'} \kappa^g(\gamma + \varepsilon \bar{\eta}, \tau)_{u,v} \sum_{i,j,k=1}^e \partial_{x^{ij} y^k}^3 g(\gamma_u + \varepsilon \bar{\eta}_u, \tau_v) \eta_u^j d(\gamma_u^i + \varepsilon \bar{\eta}_u^i) d\tau_v^k \right. \\
&\quad \left. - \int_s^{s'} \int_t^{t'} \kappa^g(\gamma, \tau)_{u,v} \sum_{i,j,k=1}^e \partial_{x^{ij} y^k}^3 g(\gamma_u, \tau_v) \eta_u^j d\gamma_u^i d\tau_v^k \right\} \\
&+ \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \left\{ \int_s^{s'} \int_t^{t'} \kappa^g(\gamma + \varepsilon \bar{\eta}, \tau)_{u,v} \sum_{i,k=1}^e \partial_{x^i y^k}^2 g(\gamma_u + \varepsilon \bar{\eta}_u, \tau_v) d\eta_u^i d\tau_v^k \right. \\
&\quad \left. - \int_s^{s'} \int_t^{t'} \kappa^g(\gamma, \tau)_{u,v} \sum_{i,k=1}^e \partial_{x^i y^k}^2 g(\gamma_u, \tau_v) d\eta_u^i d\tau_v^k \right\}.
\end{aligned}$$

Thanks to Lemma A.1, we can use dominated convergence to swap limits and integrals. After rearranging terms we obtain

$$\begin{aligned}
\partial_{\gamma}^{\eta\bar{\eta}} \kappa^g(\gamma, \tau)_{s', t'} &= \int_s^{s'} \int_t^{t'} \partial_{\gamma}^{\eta\bar{\eta}} \kappa^g(\gamma, \tau)_{u,v} \sum_{i,k=1}^e \partial_{x^i y^k}^2 g(\gamma_u, \tau_v) d\gamma_u^i d\tau_v^k \\
&+ \int_s^{s'} \int_t^{t'} \partial_{\gamma}^{\eta} \kappa^g(\gamma, \tau)_{u,v} \sum_{i,k,l=1}^e \partial_{x^{il} y^k}^3 g(\gamma_u, \tau_v) \bar{\eta}_u^l d\gamma_u^i d\tau_v^k \\
&+ \int_s^{s'} \int_t^{t'} \partial_{\gamma}^{\eta} \kappa^g(\gamma, \tau)_{u,v} \sum_{i,k=1}^e \partial_{x^i y^k}^2 g(\gamma_u, \tau_v) d\bar{\eta}_u^i d\tau_v^k \\
&+ \int_s^{s'} \int_t^{t'} \partial_{\gamma}^{\bar{\eta}} \kappa^g(\gamma, \tau)_{u,v} \sum_{i,j,k=1}^e \partial_{x^{ij} y^k}^3 g(\gamma_u, \tau_v) \eta_u^j d\gamma_u^i d\tau_v^k \\
&+ \int_s^{s'} \int_t^{t'} \kappa^g(\gamma, \tau)_{u,v} \sum_{i,j,k,l=1}^e \partial_{x^{ijl} y^k}^4 g(\gamma_u, \tau_v) \bar{\eta}_u^l \eta_u^j d\gamma_u^i d\tau_v^k \\
&+ \int_s^{s'} \int_t^{t'} \kappa^g(\gamma, \tau)_{u,v} \sum_{i,j,k=1}^e \partial_{x^{ij} y^k}^3 g(\gamma_u, \tau_v) \eta_u^j d\bar{\eta}_u^i d\tau_v^k \\
&+ \int_s^{s'} \int_t^{t'} \partial_{\gamma}^{\bar{\eta}} \kappa^g(\gamma, \tau)_{u,v} \sum_{i,k=1}^e \partial_{x^i y^k}^2 g(\gamma_u, \tau_v) d\eta_u^i d\tau_v^k \\
&+ \int_s^{s'} \int_t^{t'} \kappa^g(\gamma, \tau)_{u,v} \sum_{i,k,l=1}^e \partial_{x^{il} y^k}^3 g(\gamma_u, \tau_v) \bar{\eta}_u^l d\eta_u^i d\tau_v^k
\end{aligned}$$

where we also used the property (3.13). We can regroup the terms according to the order of κ^g , using $\sum_{i,k=1}^e \partial_{x^i y^k}^2 g(\gamma_u, \tau_v) d\gamma_u^i d\tau_v^k = \langle d\gamma_u^g, d\tau_v^g \rangle_{\mathcal{H}_g}$, Equation (A.15) and similarly for the next order:
(A.16)

$$\left\langle d\partial_{\gamma}^{\eta\bar{\eta}} G(\gamma)_u, d\tau_v^g \right\rangle_{\mathcal{H}_g} = \sum_{i,k,l=1}^e \left\{ \sum_{j=1}^e \partial_{x^i y^k}^4 g(\gamma_u, \tau_v) \bar{\eta}_u^l \eta_u^j d\gamma_u^i + \partial_{x^i y^k}^3 g(\gamma_u, \tau_v) \eta_u^j d\bar{\eta}_u^i + \partial_{x^i y^k}^3 g(\gamma_u, \tau_v) \bar{\eta}_u^l d\eta_u^i \right\} d\tau_v^k.$$

Eventually this yields the more compact formulation

$$\begin{aligned} \partial_{\gamma}^{\eta\bar{\eta}} \kappa^g(\gamma, \tau)_{s,t} &= \int_s^{s'} \int_t^{t'} \partial_{\gamma}^{\eta\bar{\eta}} \kappa^g(\gamma, \tau)_{u,v} \langle d\gamma_u^g, d\tau_v^g \rangle_{\mathcal{H}_g} + \int_s^{s'} \int_t^{t'} \partial_{\gamma}^{\eta\bar{\eta}} \kappa^g(\gamma, \tau)_{u,v} \langle d\partial_{\gamma}^{\eta\bar{\eta}} G(\gamma)_u, d\tau_v^g \rangle_{\mathcal{H}_g} \\ &\quad + \int_s^{s'} \int_t^{t'} \partial_{\gamma}^{\eta\bar{\eta}} \kappa^g(\gamma, \tau)_{u,v} \langle d\partial_{\gamma}^{\eta\bar{\eta}} G(\gamma)_u, d\tau_v^g \rangle_{\mathcal{H}_g} + \int_s^{s'} \int_t^{t'} \kappa^g(\gamma, \tau)_{u,v} \langle d\partial_{\gamma}^{\eta\bar{\eta}} G(\gamma)_u, d\tau_v^g \rangle_{\mathcal{H}_g}, \end{aligned}$$

as claimed. \square

REFERENCES

- [1] E. Abi Jaber and S. Villeneuve. Gaussian agency problems with memory and linear contracts. *Available at SSRN 4226543*, 2022.
- [2] N. Aronszajn. Theory of reproducing kernels. *Transactions of the American mathematical society*, 68(3):337–404, 1950.
- [3] O. E. Barndorff-Nielsen, M. S. Pakkanen, and J. Schmiegel. Assessing relative volatility/intermittency/energy dissipation. *Electron. J. Statist.*, 8(2):1996–2021, 2014.
- [4] O. E. Barndorff-Nielsen and J. Schmiegel. Brownian semistationary processes and volatility/intermittency. *Advanced financial modelling*, 8:1–26, 2009.
- [5] P. Battle, Y. Chen, B. Hosseini, H. Owahdi, and A. M. Stuart. Error analysis of kernel/GP methods for nonlinear and parametric PDEs. *arXiv preprint arXiv:2305.04962*, 2023.
- [6] C. Bayer, P. Friz, and J. Gatheral. Pricing under rough volatility. *Quantitative Finance*, 16(6):887–904, 2016.
- [7] C. Bayer, P. K. Friz, M. Fukasawa, J. Gatheral, A. Jacquier, and M. Rosenbaum. *Rough volatility*. SIAM, 2023.
- [8] M. Bennedsen, A. Lunde, and M. S. Pakkanen. Hybrid scheme for Brownian semistationary processes. *Finance and Stochastics*, 21:931–965, 2017.
- [9] G. Blanchard, G. Lee, and C. Scott. Generalizing from several related classification tasks to a new unlabeled sample. *Advances in neural information processing systems*, 24, 2011.
- [10] O. Bonesini, A. Jacquier, and A. Pannier. Rough volatility, path-dependent PDEs and weak rates of convergence. *arXiv preprint arXiv:2304.03042*, 2023.
- [11] T. Cass, T. Lyons, and X. Xu. Weighted signature kernels. *The Annals of Applied Probability*, 34(1A):585–626, 2024.
- [12] Y. Chen, B. Hosseini, H. Owahdi, and A. M. Stuart. Solving and learning nonlinear PDEs with gaussian processes. *arXiv preprint arXiv:2103.12959*, 2021.
- [13] Y. Chen, H. Owahdi, and F. Schäfer. Sparse Cholesky factorization for solving nonlinear PDEs via Gaussian processes. *arXiv preprint arXiv:2304.01294*, 2023.
- [14] I. Chevyrev and H. Oberhauser. Signature moments to characterize laws of stochastic processes. *arXiv preprint arXiv:1810.10971*, 2018.
- [15] A. Christmann and I. Steinwart. Consistency and robustness of kernel-based regression in convex risk minimization. *Bernoulli*, 13(3):799–819, 2007.
- [16] N. M. Cirone, M. Lemerrier, and C. Salvi. Neural signature kernels as infinite-width-depth-limits of controlled ResNets. *International Conference on Machine Learning*, 2023.
- [17] T. Cochrane, P. Foster, V. Chhabra, M. Lemerrier, T. Lyons, and C. Salvi. Sk-tree: a systematic malware detection algorithm on streaming trees via the signature kernel. In *2021 IEEE international conference on cyber security and resilience (CSR)*, pages 35–40. IEEE, 2021.
- [18] J. Cockayne, C. Oates, T. Sullivan, and M. Girolami. Probabilistic numerical methods for partial differential equations and Bayesian inverse problems. *arXiv preprint arXiv:1605.07811*, 2016.
- [19] J. M. Corcuera, E. Hedevarang, M. S. Pakkanen, and M. Podolskij. Asymptotic theory for brownian semi-stationary processes with application to turbulence. *Stochastic processes and their applications*, 123(7):2552–2574, 2013.
- [20] C. Cuchiero, P. Schmock, and J. Teichmann. Global universal approximation of functional input maps on weighted spaces. *arXiv preprint arXiv:2306.03303*, 2023.
- [21] B. Dupire and V. Tisserot-Daguette. Functional expansions. *arXiv preprint arXiv:2212.13628*, 2022.

- [22] K. Eichinger, C. Kuehn, and A. Neamtu. Sample paths estimates for stochastic fast-slow systems driven by fractional brownian motion. *Journal of Statistical Physics*, 179(5):1222–1266, 2020.
- [23] B. Fang, H. Ni, and Y. Wu. A neural rde-based model for solving path-dependent PDEs. *arXiv preprint arXiv:2306.01123*, 2023.
- [24] A. Fermanian, T. Lyons, J. Morrill, and C. Salvi. New directions in the applications of rough path theory. *IEEE BITS the Information Theory Magazine*, 2023.
- [25] P. K. Friz and N. B. Victoir. *Multidimensional stochastic processes as rough paths: theory and applications*, volume 120. Cambridge University Press, 2010.
- [26] J. Gehring and X.-M. Li. Functional limit theorems for the fractional ornstein–uhlenbeck process. *Journal of Theoretical Probability*, 35(1):426–456, 2022.
- [27] M. Hairer. Ergodicity of stochastic differential equations driven by fractional brownian motion. *Ann. Probab.*, 33(2):703–758, 2005.
- [28] J. Han, A. Jentzen, and W. E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- [29] B. Horvath, M. Lemerrier, C. Liu, T. Lyons, and C. Salvi. Optimal stopping via distribution regression: a higher rank signature approach. *arXiv preprint arXiv:2304.01479*, 2023.
- [30] Z. Issa, B. Horvath, M. Lemerrier, and C. Salvi. Non-adversarial training of neural SDEs with signature kernel scores. *Advances in Neural Information Processing Systems*, 2023.
- [31] A. Jacquier and M. Oumgari. Deep curve-dependent PDEs for affine rough volatility. *SIAM Journal on Financial Mathematics*, 14(2):353–382, 2023.
- [32] P. Kidger, P. Bonnier, I. Perez Arribas, C. Salvi, and T. Lyons. Deep signature transforms. *Advances in Neural Information Processing Systems*, 32, 2019.
- [33] F. J. Király and H. Oberhauser. Kernels for sequentially ordered data. *Journal of Machine Learning Research*, 20, 2019.
- [34] M. Lemerrier, C. Salvi, T. Damoulas, E. Bonilla, and T. Lyons. Distribution regression for sequential data. In *International Conference on Artificial Intelligence and Statistics*, pages 3754–3762. PMLR, 2021.
- [35] X.-M. Li and J. Sieber. Slow-fast systems with fractional environment and dynamics. *The Annals of Applied Probability*, 32(5):3964–4003, 2022.
- [36] T. J. Lyons, M. Caruana, and T. Lévy. *Differential equations driven by rough paths*. Springer, 2007.
- [37] G. Manten, C. Casolo, E. Ferrucci, S. W. Mogensen, C. Salvi, and N. Kilbertus. Signature kernel conditional independence tests in causal discovery for stochastic processes. *arXiv preprint arXiv:2402.18477*, 2024.
- [38] Y. Mishura, S. Ottaviano, and T. Vargiolu. Gaussian volterra processes as models of electricity markets. *arXiv preprint arXiv:2311.09384*, 2023.
- [39] J. Morrill, C. Salvi, P. Kidger, and J. Foster. Neural rough differential equations for long time series. In *International Conference on Machine Learning*, pages 7829–7838. PMLR, 2021.
- [40] A. Pannier. Path-dependent PDEs for volatility derivatives. *arXiv preprint arXiv:2311.08289*, 2023.
- [41] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.
- [42] Z. Ren and X. Tan. On the convergence of monotone schemes for path-dependent PDEs. *Stochastic Processes and their Applications*, 127(6):1738–1762, 2017.
- [43] M. Sabate-Vidales, D. Šiška, and L. Szpruch. Solving path dependent PDEs with LSTM networks and path signatures. *arXiv preprint arXiv:2011.10630*, 2020.
- [44] C. Salvi, T. Cass, J. Foster, T. Lyons, and W. Y. The signature kernel is the solution of a Goursat PDE. *SIAM Journal on Mathematics of Data Science*, 3(3):873–899, 2021.
- [45] C. Salvi, M. Lemerrier, T. Cass, E. V. Bonilla, T. Damoulas, and T. J. Lyons. SigGPDE: Scaling sparse Gaussian processes on sequential data. In *International Conference on Machine Learning*, pages 6233–6242. PMLR, 2021.
- [46] C. Salvi, M. Lemerrier, and A. Gerasimovics. Neural stochastic PDEs: Resolution-invariant learning of continuous spatiotemporal dynamics. *Advances in Neural Information Processing Systems*, 35:1333–1344, 2022.
- [47] C. Salvi, M. Lemerrier, C. Liu, B. Horvath, T. Damoulas, and T. Lyons. Higher order kernel mean embeddings to capture filtrations of stochastic processes. *Advances in Neural Information Processing Systems*, 34:16635–16647, 2021.
- [48] Y. F. Saporito and Z. Zhang. Path-dependent deep Galerkin method: A neural network approach to solve path-dependent partial differential equations. *SIAM Journal on Financial Mathematics*, 12(3):912–940, 2021.
- [49] J. Sirignano and K. Spiliopoulos. DGM: A deep learning algorithm for solving partial differential equations. *Journal of computational physics*, 375:1339–1364, 2018.
- [50] D. Sonechkin. Climate dynamics as a nonlinear brownian motion. *International Journal of Bifurcation and Chaos*, 8(04):799–803, 1998.

- [51] B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. R. Lanckriet. Hilbert space embeddings and metrics on probability measures. *The Journal of Machine Learning Research*, 11:1517–1561, 2010.
- [52] F. Viens and J. Zhang. A martingale approach for fractional Brownian motions and related path dependent PDEs. *The Annals of Applied Probability*, 29(6):3489–3540, 2019.
- [53] H. Wang, J. Yong, and J. Zhang. Path dependent Feynman–Kac formula for forward backward stochastic Volterra integral equations. *Annales de l’Institut Henri Poincaré: Probabilités et Statistiques*, 58(2):603–638, 2022.
- [54] H. Wendland. *Scattered data approximation*, volume 17. Cambridge university press, 2004.
- [55] D.-X. Zhou. Derivative reproducing properties for kernel methods in learning theory. *Journal of computational and Applied Mathematics*, 220(1-2):456–463, 2008.