OPTIMAL STOPPING WITH SIGNATURES

C. BAYER, P. HAGER, S. RIEDEL, AND J. SCHOENMAKERS

ABSTRACT. We propose a new method for solving optimal stopping problems (such as American option pricing in finance) under minimal assumptions on the underlying stochastic process X. We consider classic and randomized stopping times represented by linear and non-linear functionals of the *rough path signature* $\mathbb{X}^{<\infty}$ associated to X, and prove that maximizing over these classes of *signature stopping times*, in fact, solves the original optimal stopping problem. Using the algebraic properties of the signature, we can then recast the problem as a (deterministic) optimization problem depending only on the (truncated) expected signature $\mathbb{E}\left[\mathbb{X}_{0,T}^{\leq N}\right]$. By applying a deep neural network approach to approximate the non-linear signature functionals, we can efficiently solve the optimal stopping problem numerically. The only assumption on the process X is that it is a continuous (geometric) random rough path. Hence, the theory encompasses processes such as fractional Brownian motion, which fail to be either semi-martingales or Markov processes, and can be used, in particular, for American-type option pricing in fractional models, e.g. on financial or electricity markets.

CONTENTS

1. Introduction

The theory of *rough paths*, see, for instance, [LCL07, FV10, FH14], provides a powerful and elegant pathwise theory of stochastic differential equations driven by general classes of stochastic processes – or, more precisely, rough paths. One of the benefits of the theory is that the resulting solution maps are continuous rather than merely measurable as in the Itō version. This property has lead to many important theoretical progresses, most notably perhaps Hairer's theory for singular non-linear SPDEs.

In addition to the theoretical advances, tools from rough path analysis – specifically, the path *signature* – play an increasingly prominent role in applications, most notably in *machine learning*, see, e.g., [AGG+18]. Intuitively, the signature $\mathbb{X}^{<\infty}$ of a path $X : [0, T] \to \mathbb{R}^d$ denotes the (infinite) collection of all iterated integrals of all components of the path against each other, i.e., of the form

$$\int_{0<t_1<\cdots<t_n<T} dX_{t_1}^{i_1} \cdots dX_{t_n}^{i_n},$$

$i_1, \ldots, i_n \in \{1, \ldots, d\}$, $n \geq 0$. To better understand the importance of the signature, let us first recall that the signature $\mathbb{X}^{<\infty}$ determines the underlying path X (up to "tree-like excursions"), which was first proved in [HL10] for paths X of bounded variation and later extended to less regular paths. This implies that, in principle, we can always work with the signature rather than the path. (A somewhat dubious proposition, as we merely replace one infinite dimensional object by another one.) However, the signature is not an arbitrary encoding of the path. Rather, Lyons' universal limit theorem suggests that the solution of a differential equation driven by a rough path X can be approximated with high accuracy by relatively few terms of the signature $\mathbb{X}^{<\infty}$. In that sense, an appropriately truncated signature can be seen as a highly efficient *compression* of X, at least in the context of dynamical systems. And, indeed, there is now ample evidence of the power of the signature as a *feature* in the sense of machine learning.

This paper is motivated by another recent application of signatures, namely the solution of stochastic optimal control problems in finance. We follow the presentation of [KLPA20], where a signature-based approach for solving optimal execution problems is developed. In a nutshell, the strategy can be summarized as follows:

(1) Trading strategies for execution of a position can be understood as (continuous) functionals $\phi(X|_{[0,t]})$ of the price path, and, hence, as functionals $\theta(\mathbb{X}_{0,t}^{<\infty})$ of the signature (at least approximately).

(2) Taking advantage of the algebraic structure of the signature (see Section 2 below), we may efficiently approximate continuous functionals $\theta(\mathbb{X}_{0,t}^{<\infty})$ by *linear* functionals $\langle l, \mathbb{X}_{0,t}^{<\infty} \rangle$, which further extends to the whole value function.

(3) Taking advantage of the linearity, we may interchange the expectation with the linear functional, thereby reducing the optimal control problem to a problem of maximizing $l \mapsto \langle l, \mathbb{E}[\mathbb{X}_{0,t}^{<\infty}] \rangle$ over some set of dual elements l.

(4) Truncate the expected signature to a finite level N.

The above strategy, in principle, only imposes very mild conditions on the underlying process X, mainly that it is continuous but possibly rough. In particular, X does not need to be a Markov

process or a semi-martingale. For this reason, we may consider the approach to be model-free.[1] Note, however, that the assumption of the existence of the expected signature $\mathbb{E}[\mathbb{X}_{0,t}^{<\infty}]$ is a rather strong assumption – in particular, ruling out many stochastic volatility models, such as the Heston model.

[KLPA20] contains extensive numerical examples, indicating the method's excellent performance in various scenarios and models, often beating benchmark methods from the financial literature. On the other hand, theoretical justification of the different approximation steps summarized above is largely missing.

We extend the method of [KLPA20] to another important control problem in finance, namely the *optimal stopping* problem, or, in more financial terms, the pricing of *American options*, for which we provide rigorous theory. More specifically, we are concerned with the problem of computing

$$\tag{1.1} \sup_{\tau \in \mathcal{S}} \mathbb{E}\left[Y_{\tau \wedge T}\right],$$

where Y denotes a process adapted to the filtration $(\mathcal{F}_t)_{t \in [0,T]}$ generated by a rough path process $(X_t)_{t \in [0,T]}$, and $\mathcal{S}$ denotes the set of all stopping times w.r.t. the same filtration. In a financial context, Y usually denotes a reward process discounted with respect to some numéraire. At first glance the optimal stopping problem (1.1) may seem unsuitable for the signature-based approach, as typical candidate stopping times are hitting times of sets, which are generally discontinuous w.r.t. the underlying path. We solve this issue by using *randomized stopping times*, see [BTW20]. Note that extending the set $\mathcal{S}$ to also include randomized stopping times does not change the value of (1.1). In the end, we are able to prove that replacing proper stopping times by *signature stopping times* – i.e., stopping times given in terms of linear functionals of the signature $\mathbb{X}^{<\infty}$ – does not change the value of the optimal stopping problem either. More precisely, we have

Theorem 1.1. *Assume that* $\mathbb{E}[\|Y\|_\infty] < \infty$. *Then,*

$$\sup_{\tau \in \mathcal{S}} \mathbb{E}[Y_{\tau \wedge T}] = \sup_{\tau_l} \mathbb{E}[Y_{\tau_l \wedge T}],$$

where the supremum on the right-hand-side ranges over stopping times $\tau_l := \inf\{t \in [0,T] : \langle l, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \geq 1\}$ *connected with linear functionals l on the signature process $\widehat{\mathbb{X}}_{0,t}^{<\infty}$ (we refer to Sections 2 and 4 for precise definitions).*

The theorem is presented as Proposition 5.5 below. We note that, following [KLPA20], we extend the path X by adding running time as an additional component. $\widehat{\mathbb{X}}^{<\infty}$ denotes the signature of the extended path.

In the next step we need to actually compute a maximizing signature stopping time. In this context, this most importantly implies replacing the full signature $\widehat{\mathbb{X}}^{<\infty}$ by a truncated version $\widehat{\mathbb{X}}^{\leq N}$. Using some further technical assumptions, Proposition 6.5 provides convergence of the corresponding approximations to the value of (1.1) as $N \to \infty$.

Assuming that Y is a polynomial function of X – or, more generally, of $\mathbb{X}^{<\infty}$ – we can derive an approximation formula in terms of an optimization problem involving linear functionals of the

---

[1]The full expected signature $\mathbb{E}[\mathbb{X}^{<\infty}]$ typically characterizes the distribution of the process X, see [CL16]. In that sense, we would hesitate to regard methods relying on the full (rather than truncated) expected signature as "model-free".

expected signature $\mathbb{E}[\widehat{\mathbb{X}}^{\leq N}]$ rather than the expectation of some functional of the signature. See Corollary 6.6 for details.

While this approach undoubtedly has many attractive features, very deep truncation levels N might be needed for complicated optimal stopping problems. Hence, we also consider stopping rules which are parameterized by non-linear functionals of the (truncated) signature. More precisely, we consider stopping rules represented by *deep neural networks* (DNNs) applied to the *log-signature*. DNNs are well known for their combination of high expressiveness, especially in high dimensions, relative ease of training, and wide availability of efficient software implementations. On the other hand, the log-signature removes linear redundancies from the full signature, and, therefore, provides a compressed but equivalent representation of the data contained in the signature. We prove the following analogue to Theorem 1.1, see Proposition 7.4 for details of the construction and the proof.

Theorem 1.2. *Let $\mathcal{T}_{\log}$ denote the set of DNNs θ of a suitable chosen architecture, chosen such that θ accepts a truncated log-signature $\log^{\otimes} \widehat{\mathbb{X}}_{0,s}^{\leq N}$ truncated at some level $N = N(\theta)$, $N = 1, 2, \ldots$. Define the* deep, randomized signature stopping rule

$$\tau_\theta^r := \inf \left\{ t \geq 0 \; : \; \int_0^t \theta \left(\log^{\otimes} \widehat{\mathbb{X}}_{0,s}^{\leq N} \right)^2 ds \geq Z \right\},$$

for an independent r.v. $Z \sim \mathrm{Exp}(1)$. Assuming that $\mathbb{E}\left[\|Y\|_\infty\right] < \infty$, we then have

$$\sup_{\tau \in \mathcal{S}} \mathbb{E}\left[Y_{\tau \wedge T}\right] = \sup_{\theta \in \mathcal{T}_{\log}} \mathbb{E}\left[Y_{\tau_\theta^r \wedge T}\right].$$

Finally, Section 8 contains extensive numerical examples of both the linearized and the deep log signature method. We consider two models. Motivated by [BCJ19] we consider the problem of optimally stopping a fractional Brownian motion. In this example, we find excellent performance of the deep signature stopping approach. En passant, we provide explicit formulas for the limiting case $H \to 0$. We also consider an example from optimal control in an electricity market with the price function modeled by a fractional Brownian motion, as motivated by [Ben17].

Outline of the paper. Section 2 recalls basic definitions from the theory of rough paths and provides the algebraic and analytic setting of signatures. A framework for studying stopped rough paths is presented in Section 3. Stopping times based on continuous functionals on rough paths and their randomized counter-parts are introduced and discussed in Section 4. The special case of linear functionals of the signature is studied in Section 5. The following Section 6 contains a fully linear approximation to the optimal stopping problems in terms of the expected truncated signature. Deep signature stopping rules are introduced in Section 7. Finally, numerical examples are presented in Section 8.

2. Preliminaries

We start by introducing the basic definitions needed for understanding signatures and their algebraic and – in the context of rough paths – analytic properties. These definitions are standard in the rough path literature, we refer to [LCL07, FH14, FV10] for a more detailed exposition.

2.1. **The tensor algebra.** Let V be a finite-dimensional $\mathbb{R}$-vector space with basis $\{e_1, \ldots, e_d\}$. The dual space is denoted by V^* with dual basis $\{e_1^*, \ldots, e_d^*\}$. We define the *tensor algebra* and the *extended tensor algebra* by setting

$$T(V) := \bigoplus_{n=0}^{\infty} V^{\otimes n} \quad \text{and} \quad T((V)) := \prod_{n=0}^{\infty} V^{\otimes n}$$

where $V^{\otimes n}$ denotes the n-th tensor power of V with the convention $V^{\otimes 0} := \mathbb{R}$, $V^{\otimes 1} := V$. Note that there is a natural pairing between $T((V))$ and $T(V^*)$ which we denote by

$$\langle \cdot, \cdot \rangle \colon T(V^*) \times T((V)) \to \mathbb{R}.$$

We define sum and product of two elements $\mathbf{a} = (a_n)_{n=0}^{\infty}, \mathbf{b} = (b_n)_{n=0}^{\infty} \in T((V))$ by setting

$$\mathbf{a} + \mathbf{b} := (a_n + b_n)_{n=0}^{\infty},$$

$$\mathbf{a} \otimes \mathbf{b} := \Big(\sum_{i+j=n} a_i \otimes b_j \Big)_{n=0}^{\infty}.$$

For $\lambda \in \mathbb{R}$, we define $\lambda \mathbf{a} := (\lambda a_n)_{n=0}^{\infty}$. We also let $\mathbf{0} := (0, 0, \ldots)$ and $\mathbf{1} := (1, 0, 0, \ldots)$. Note that

$$\mathbf{1} \otimes \mathbf{a} = \mathbf{a} \otimes \mathbf{1} = \mathbf{a}$$

for every $\mathbf{a} \in T((V))$. The *truncated tensor algebra* is defined by

$$T^N(V) := \bigoplus_{n=0}^{N} V^{\otimes n}.$$

We define maps $\pi_n \colon T((V)) \to V^{\otimes n}$ and $\pi_{\leq N} \colon T((V)) \to T^N(V)$ by $\pi_n(\mathbf{a}) = a_n$ and $\pi_{\leq N}(\mathbf{a}) = (a_0, \ldots, a_N)$ where $\mathbf{a} = (a_n)_{n=0}^{\infty}$. We will sometimes abuse notation and write $\mathbf{0}$ and $\mathbf{1}$ for the elements $\pi_{\leq N}(\mathbf{0})$ and $\pi_{\leq N}(\mathbf{1})$ in the truncated tensor algebra.

Next, we consider norms on $T((V))$ and $T(V^*)$. On V, we choose the l^{∞}-norm, i.e. for $v = \lambda_1 e_1 + \ldots + \lambda_d e_d$, we set $|v| := \max_i |\lambda_i|$. For elements in V^*, we use the l^1-norm, i.e. $|v^*| := |\lambda_1| + \ldots + |\lambda_d|$ for $v^* = \lambda_1 e_1^* + \ldots + \lambda_d e_d^*$. On the tensor powers of V resp. V^*, we use the corresponding norms, too. Note that the norms on the tensor products $V^{\otimes n}$ are *admissible*, meaning that if $v = a_1 \otimes \ldots \otimes a_k$ and $\sigma v := a_{\sigma(1)} \otimes \ldots \otimes a_{\sigma(k)}$ for a permutation σ, $|\sigma v| = |v|$, and $|v \otimes w| \leq |v||w|$. We set

$$|\mathbf{a}| := \sup_{i \in \mathbb{N}_0} |a_i| \in [0, \infty] \quad \text{for } \mathbf{a} = (a_i)_{i=0}^{\infty} \in T((V))$$

and

$$|\mathbf{b}| := \sum_{i=0}^{\infty} |b_i| \in [0, \infty) \quad \text{for } \mathbf{b} = (b_i)_{i=0}^{\infty} \in T(V^*).$$

Note that we always have

$$|\langle \mathbf{b}, \mathbf{a} \rangle| \leq |\mathbf{b}||\pi_{\leq N}(\mathbf{a})| \leq |\mathbf{b}||\mathbf{a}|$$

where $N = \max\{i \in \mathbb{N}_0 \,:\, b_i \neq 0\}$.

2.2. **Shuffles.** In the following, calculations will mainly be performed in the space $T(V^*)$. In order to simplify notations, we will replace expressions like $e_{i_1}^* \otimes \cdots \otimes e_{i_n}^*$ by the much simpler form $i_1 \cdots i_n$. More precisely, let $\mathcal{W}(\mathcal{A}_d)$ be the linear span of words composed by the letters in the dictionary $\mathcal{A}_d = \{1, \ldots, d\}$. The empty word is denoted by $\varnothing \in \mathcal{W}(\mathcal{A}_d)$. We can naturally define the sum $l_1 + l_2$ and the scalar product λl for elements $l, l_1, l_2 \in \mathcal{W}(\mathcal{A}_d)$ and $\lambda \in \mathbb{R}$. If $w = i_1 \cdots i_n$ and $v = j_1 \cdots j_m$ are two words, the *concatenation* is defined by

$$wv := i_1 \cdots i_n j_1 \cdots j_m.$$

This operation is extended bi-linearly to elements in $\mathcal{W}(\mathcal{A}_d)$. The basis elements $\{e_{i_1}^* \otimes \cdots \otimes e_{i_n}^* : i_1, \ldots, i_n \in \{1, \ldots, d\}\}$ in $(V^*)^{\otimes n}$ can be identified with words via the map

$$e_{i_1}^* \otimes \cdots \otimes e_{i_n}^* \mapsto i_1 \cdots i_n$$

which induces an isomorphism $T(V^*) \cong \mathcal{W}(\mathcal{A}_d)$. We can also think of $\mathcal{W}(\mathcal{A}_d)$ as the space of non-commutative polynomials where the unknown are given by the letters $\{1, \ldots, d\}$. For a word $w = i_1 \cdots i_n$, set $\deg(w) := n$ and $\deg(\varnothing) := 0$. If $l = \lambda_1 w_1 + \ldots + \lambda_n w_n \in \mathcal{W}(\mathcal{A}_d)$ with $\lambda_1, \ldots, \lambda_n \in \mathbb{R} \setminus \{0\}$ and $w_1, \ldots, w_n$ words, we define

$$\deg(l) := \max_{i=1 \ldots, n} \deg(w_i).$$

Apart from concatenation, there is a second important product defined on $\mathcal{W}(\mathcal{A}_d)$ which is called *shuffle product*: For a word w, we set

$$w \sqcup \varnothing := \varnothing \sqcup w := w.$$

If wi and vj are words and $i, j \in \mathcal{A}_d$ are letters, we recursively define $wi \sqcup vj \in \mathcal{W}(\mathcal{A}_d)$ by

$$wi \sqcup vj := (w \sqcup vj)i + (wi \sqcup v)j.$$

This operation is extended bi-linearly to a product $\sqcup : \mathcal{W}(\mathcal{A}_d) \times \mathcal{W}(\mathcal{A}_d) \to \mathcal{W}(\mathcal{A}_d)$. Note that $\sqcup$ is associative, commutative and distributive over $+$. If $P \in \mathbb{R}[x]$ is a commutative polynomial with unknown variable x, i.e. $P(x) = \lambda_0 + \lambda_1 x + \ldots \lambda_n x^n$, we define $P^{\sqcup} : \mathcal{W}(\mathcal{A}_d) \to \mathcal{W}(\mathcal{A}_d)$ by setting

$$(2.1) \qquad P^{\sqcup}(l) := \lambda_0 \varnothing + \lambda_1 l + \lambda_2 (l \sqcup l) + \ldots + \lambda_n l^{\sqcup n},$$

where $l^{\sqcup k}$ denotes k-th shuffle product of $l \in \mathcal{W}(\mathcal{A}_d)$ with itself.

We define

$$(2.2) \qquad G(V) := \{\mathbf{a} \in T((V)) \setminus \{\mathbf{0}\} : \langle l_1 \sqcup l_2, \mathbf{a}\rangle = \langle l_1, \mathbf{a}\rangle\langle l_2, \mathbf{a}\rangle \text{ for every } l_1, l_2 \in T(V^*)\}$$

and call it the set of *group-like elements*. Note that $\pi_0(\mathbf{g}) = 1$ for every $\mathbf{g} \in G(V)$. One can show that $(G(V), \otimes)$ is a group with identity $\mathbf{1}$ and inverse given by

$$\mathbf{g}^{-1} = \sum_{n \geq 0} (\mathbf{1} - \mathbf{g})^{\otimes n}.$$

We also set $G^N(V) := \pi_{\leq N}(G(V))$ which is a free nilpotent group of order N with respect to the "truncated multiplication" $\mathbf{a} \otimes_{G^N(V)} \mathbf{b} := \pi_N(\mathbf{a} \otimes \mathbf{b})$, for $\mathbf{a}, \mathbf{b} \in G^N(V)$. However, we will not distinguish between the multiplication symbols on $G^N(V)$ and $G(V)$ and use $\otimes$ in both cases.

Remark 2.1. The relation $\langle l_1 \sqcup l_2, \mathbf{a}\rangle = \langle l_1, \mathbf{a}\rangle\langle l_2, \mathbf{a}\rangle$ for $\mathbf{a} \in G(V)$ implies that

$$(2.3) \qquad P(\langle l, \mathbf{a}\rangle) = \langle P^{\sqcup}(l), \mathbf{a}\rangle$$

for any polynomial P. This is really the justification for introducing the shuffle product, as it provides an explicit linearization of polynomials in the signature.

2.3. **Rough paths and their signatures.** Now that the algebraic setting for signatures is developed (for the purposes of this paper), we can finally consider the analytic properties of (rough) paths. More concretely, given a path $X : [0, T] \to V$ (of sufficient regularity), we will associate to it a function $\mathbb{X}$ taking values in the truncated tensor algebra, which is the fundamental building block of rough path theory. Set $\Delta_T := \{(s, t) \in [0, T]^2 : 0 \le s \le t \le T\}$. For a map $\mathbb{X} \colon \Delta_T \to T^N(V)$, we define its *p-variation*

$$\|\mathbb{X}\|_{p-\mathrm{var};[s,t]} := \max_{k=1,\ldots,N} \sup_{\mathcal{D} \subset [s,t]} \left(\sum_{t_i \in \mathcal{D}} |\pi_k(\mathbb{X}_{t_i, t_{i+1}})|^{\frac{p}{k}} \right)^{\frac{k}{p}}$$

where the supremum ranges over all partitions $\mathcal{D}$ of $[s, t]$. We will use the notation $\|\mathbb{X}\|_{p-\mathrm{var}} := \|\mathbb{X}\|_{p-\mathrm{var};[0,T]}$. For $\mathbb{X}, \mathbb{Y} \colon \Delta_T \to T^N(V)$, we define the *p-variation distance*

$$d_{p-\mathrm{var};[s,t]}(\mathbb{X}, \mathbb{Y}) := \|\mathbb{X} - \mathbb{Y}\|_{p-\mathrm{var};[s,t]}$$

and set $d_{p-\mathrm{var}}(\mathbb{X}, \mathbb{Y}) := d_{p-\mathrm{var};[0,T]}(\mathbb{X}, \mathbb{Y})$. A *weakly geometric p-rough path* $\mathbb{X}$ is a continuous path $\mathbb{X} \colon [0, T] \to G^{\lfloor p \rfloor}(V)$ with $\mathbb{X}_0 = \mathbf{1}$ and $\|\mathbb{X}\|_{p-\mathrm{var}} < \infty$ where we set $\mathbb{X}_{s,t} := \mathbb{X}_s^{-1} \otimes \mathbb{X}_t$ for $s \le t$. Note that $\mathbb{X}_t = \mathbb{X}_{0,t}$. We denote the space of weakly geometric p-rough paths by $\mathcal{W}\Omega_T^p$ and equip it with the distance $d_{p-\mathrm{var}}$. If $X \colon [0, T] \to V$ is a continuous path of bounded variation, we define its *signature* $\mathbb{X}^{<\infty} \colon [0, T] \to T((V))$ by

$$\pi_k(\mathbb{X}_t^{<\infty}) := \int_{0 < t_1 < \ldots < t_k < t} dX_{t_1} \otimes \cdots \otimes dX_{t_k}.$$

The *truncated signature* $\mathbb{X}^{\le N} \colon [0, T] \to T^N(V)$ is defined by $\mathbb{X}^{\le N} := \pi_{\le N}(\mathbb{X}^{<\infty})$. It can be checked that $\mathbb{X}^{<\infty}$ takes values in $G(V)$ and we set $\mathbb{X}_{s,t}^{<\infty} := (\mathbb{X}_s^{<\infty})^{-1} \otimes \mathbb{X}_t^{<\infty}$ so that

$$\pi_k(\mathbb{X}_{s,t}^{<\infty}) = \int_{s < t_1 < \ldots < t_k < t} dX_{t_1} \otimes \cdots \otimes dX_{t_k}.$$

One can also show that $\mathbb{X}^{\le N}$ is an element in $\mathcal{W}\Omega_T^p$ for every $p \ge 1$ with $N = \lfloor p \rfloor$.

A *geometric p-rough path* $\mathbb{X}$ is a weakly geometric rough path $\mathbb{X} \in \mathcal{W}\Omega_T^p$ for which there exists a sequence of piecewise smooth paths (X_n) such that $d_{p-\mathrm{var}}(\mathbb{X}, \mathbb{X}_n^{\le \lfloor p \rfloor}) \to 0$ as $n \to \infty$. The space of geometric rough paths is denoted by Ω_T^p. It can be shown that the inclusion $\Omega_T^p \subset \mathcal{W}\Omega_T^p$ is strict and that Ω_T^p is a Polish space. From Lyons' Extension theorem [LCL07, Theorem 3.7], every geometric rough path $\mathbb{X} \in \Omega_T^p$ has a unique lift $\mathbb{X}^{<\infty}$ which is a path in $G(V)$, satisfying $\|\pi_{\le N}(\mathbb{X}^{<\infty})\|_{p-\mathrm{var}} < \infty$ for every $N \ge 1$ and $\pi_{\le \lfloor p \rfloor}(\mathbb{X}^{<\infty}) = \mathbb{X}$. We call $\mathbb{X}^{<\infty}$ the *signature of the rough path* $\mathbb{X}$ and $\mathbb{X}^{\le N} := \pi_{\le N}(\mathbb{X}^{<\infty})$ its *truncated signature*.

Similarly, for $V = \mathbb{R}^{1+d}$, we define the space $\widehat{\Omega}_T^p$ as the closure of rough path lifts $\widehat{\mathbb{X}}^{\le \lfloor p \rfloor}$ in the p-variation distance where $\widehat{X}_t = (t, X_t) \in \mathbb{R}^{1+d}$ and X is piecewise smooth. It follows that $\widehat{\Omega}_T^p$ is Polish.

Remark 2.2. Following the notation introduced above, the letter 1 corresponds to the running time component t of the path $\widehat{X}$, whereas the components of X correspond to the letters $2, \ldots, d+1$, respectively.

Example 2.3 (Brownian motion as a rough path). Let X be a d-dimensional Brownian motion. In this case a natural lift to a geometric rough path $\mathbb{X} \in \Omega_T^p$ with $p \in (2, 3)$ is given by

$$\mathbb{X}_{s,t} = \left(1, X_{s,t}, \int_s^t X_{s,u} \otimes \circ dX_u\right), \quad 0 \le s \le t \le T$$

where $X_{s,t} = X_t - X_s$ and for all $i, j \in \mathcal{A}_d$ the tensor valued Stratonovich integral is given by

$$\left\langle ij, \int_s^t X_{s,u} \otimes \circ dX_u \right\rangle = \int_s^t X_{s,u}^i \circ dX_u^j = \int_s^t X_{s,u}^i dX_u^j + \frac{1}{2}[X^i, X^j]_{s,t}.$$

Indeed, to see that $\mathbb{X} \in \mathcal{W}\Omega_T^p$, one may readily check that $\mathbb{X}_{s,t} \in G^2(V)$ is an immediate consequence of the product rule and the rough path regularity of $\mathbb{X}$ follows from a generalized Kolmogorov criterion (see [FH14, Theorem 3.1]). As it is well known that the integral with respect to the piecewise linear approximation of Brownian motion converges to the Stratonovich integral a.s., we also see that $\mathbb{X} \in \Omega_T^p$. The signature $\mathbb{X}^{<\infty}$ of the enhanced Brownian motion $\mathbb{X}$ is then given by the iterated integrals of all orders, i.e. for any word $w = i_1 \cdots i_k \in \mathcal{W}(\mathcal{A}_d)$ we have

$$\langle i_1 \cdots i_k, \mathbb{X}_{0,t}^{<\infty}\rangle = \int_{0<t_1<\ldots<t_k<t} \circ dX_{t_1}^{i_1} \cdots \circ dX_{t_k}^{i_k}.$$

An explicit form of the expected signature is also known due to Fawcett [Faw02]

$$\mathbb{E}(\mathbb{X}_{0,t}^{<\infty}) = \exp^{\otimes}\left(\frac{1}{2}t\sum_{i=1}^d e_i \otimes e_i\right) = \sum_{n=0}^{\infty} \frac{1}{n!}\frac{t^n}{2^n}\left(\sum_{i=1}^d e_i \otimes e_i\right)^{\otimes n}.$$

Note that this construction of a geometric rough path works in principle for all continuous semi-martingales and we refer to [FV10, Section 14] for more detail.

Example 2.4 (Fractional Brownian motion). Let X be a one-dimensional fractional Brownian motion with Hurst parameter $H \in (0, 1)$, i.e. X is a zero mean Gaussian process with covariance function

$$(2.4) \qquad \mathbb{E}(X_s X_t) = \frac{1}{2}\left(|s|^{2H} + |t|^{2H} - |t-s|^{2H}\right), \quad 0 \le s \le t.$$

Recall that the sample paths of X are a.s. $(H-\varepsilon)$-Hölder continuous for any $\varepsilon > 0$. In case $H = 1/2$ X is just a standard Brownian motion, in case $H \neq 1/2$, however, X is not a Markov process and not a semimartingale. However, since X is one-dimensional ($V = \mathbb{R}$) there is a trivial lift to a geometric rough path $\mathbb{X} \in \Omega_T^p$ for any $p \in (1/H, 1 + 1/H)$ given by

$$\mathbb{X}_{s,t} = \left(1, X_{s,t}, \frac{1}{2}(X_{s,t})^{\otimes 2}, ..., \frac{1}{\lfloor p \rfloor!}(X_{s,t})^{\otimes \lfloor p \rfloor}\right) \equiv \exp_{\lfloor p \rfloor}^{\otimes}(X_{s,t}) \in G^{\lfloor p \rfloor}(V), \quad 0 \le s \le t \le T.$$

As we will see in the next section, we are particulary interested in the process $\widehat{X}$ defined by $\widehat{X}_t = (t, X_t) \in \mathbb{R}^2$. The first component of $\widehat{X}$ is of locally bounded variation and therefore it can be lifted to a geometric rough path $\widehat{\mathbb{X}} \in \widehat{\Omega}_T^p$ (see [FV10, Theorem 9.26]). Intuitively speaking we can make use of the abundant regularity of the first component $\widehat{X}$ in order to define iterated integrals by imposing the integration by parts rule. More precisely in case $p > 2$ we have

$$\langle 12, \widehat{\mathbb{X}}_{s,t}\rangle = \langle 2, \widehat{\mathbb{X}}_{s,t}\rangle\langle 1, \widehat{\mathbb{X}}_{s,t}\rangle - \langle 21, \widehat{\mathbb{X}}_{s,t}\rangle = X_{s,t}(t-s) - \int_s^t X_{s,u}du,$$

and the right hand side is clearly well defined. Using the shuffle identity, this reasoning can be carried on to express all components of the signature $\widehat{\mathbb{X}}^{<\infty}$ in terms of increments of X, finite-variation integrals and products thereof.

3. The space of stopped rough paths

We will now consider rough paths $\mathbb{Z}$ defined on some intervals $[0, s] \subset [0, T]$. In order to naturally model the notion of *adaptedness* to a filtration, we will consider functionals of the restriction of a rough path $\mathbb{Z}$ to a subinterval of its domain. Hence, the analysis of the corresponding control problem requires us to define a distance of rough paths with different domains. Following [KLPA20], we will use a distance motivated by Dupire's functional Itō calculus, see [Dup19, CF10]. This means, when we compare a path Z^1 defined on $[0, s]$ and another path Z^2 defined on $[0, t]$ with $s < t$, we will extend Z^1 to $[0, t]$ by $Z_u^1 := Z_s^1$ for $s \le u \le t$. We will, in principle, use the same construction for rough paths, but recall that we are considering paths $u \mapsto (u, X_u)$ in our framework, and extending the time component of such a path in a constant way does not make much sense. Instead, we will apply Dupire's extension to the X-component, but use the linear extension (i.e., the exact one) for the time component.

More precisely, let $\mathbb{Z}|_{[0,s]} \in \widehat{\Omega}_s^p$ and $s \le t$. By definition, there exists a sequence $Z_u^n = (u, X_u^n)$ where $X^n \colon [0, s] \to \mathbb{R}^d$ is a piecewise smooth path such that $d_{p-\mathrm{var};[0,s]}(\mathbb{Z}|_{[0,s]}, \mathbb{Z}^{n;\le \lfloor p \rfloor}) \to 0$ as $n \to \infty$. Set $\widetilde{X}_u^n := X_{u \wedge s}^n$ for $u \in [0, t]$ and $\widetilde{Z}_u^n := (u, \widetilde{X}_u^n)$. One can check that $\widetilde{\mathbb{Z}}^{n;\le \lfloor p \rfloor}$ is a Cauchy sequence in $\widehat{\Omega}_t^p$, and we denote the limit by $\widetilde{\mathbb{Z}}|_{[0,t]}$. One can also check that the definition of $\widetilde{\mathbb{Z}}|_{[0,t]}$ does not depend on the choice of the sequence X^n. By construction we have that $\widetilde{\mathbb{Z}}|_{[0,s]} = \mathbb{Z}|_{[0,s]}$, which motivates the following definition.

Definition 3.1. For $T > 0$, we set $\Lambda_T := \bigcup_{t \in [0,T]} \widehat{\Omega}_t^p$ and call it the *space of stopped rough paths*. We equip it with the metric

$$d(\mathbb{X}|_{[0,t]}, \mathbb{Y}|_{[0,s]}) := d_{p-\mathrm{var};[0,t]}(\mathbb{X}|_{[0,t]}, \widetilde{\mathbb{Y}}|_{[0,t]}) + |t - s|$$

where we assume $s \le t$ and $\widetilde{\mathbb{Y}}|_{[0,t]}$ is the stopped rough path constructed as explained above.

Let us mention that Λ_T is Polish. For this and related simple technical facts about the topology of Λ_T, we refer to the Appendix A. Later on we will use that $\mathbb{1}_{\{\tau(\omega) \le t\}}$ can be represented as a measurable map of the restricted rough path.

Lemma 3.2. *Let $\widehat{\mathbb{X}}$ be a stochastic process in $\widehat{\Omega}_T^p$ and set $\mathcal{F}_t := \sigma(\widehat{\mathbb{X}}_{0,s} : 0 \le s \le t) = \sigma(\widehat{\mathbb{X}}|_{[0,t]})$. Let τ be a stopping time with respect to $(\mathcal{F}_t)$. Then there is a Borel measurable map $\theta \colon \Lambda_T \to \{0, 1\}$ such that*

$$\theta(\widehat{\mathbb{X}}(\omega)|_{[0,t]}) = \mathbb{1}_{\{\tau(\omega) \le t\}}$$

for every $\omega \in \Omega$.

Proof. For every $t \in [0, T]$, $\{\tau \le t\}$ is $\sigma(\widehat{\mathbb{X}}|_{[0,t]})$-measurable, hence there is a set $A_t \in \mathcal{B}(\widehat{\Omega}_t^p)$ such that $(\widehat{\mathbb{X}}|_{[0,t]})^{-1}(A_t) = \{\tau \le t\}$. It follows that

$$\tag{3.1} \mathbb{1}_{\{\tau(\omega) \le t\}} = \mathbb{1}_{A_t}(\widehat{\mathbb{X}}(\omega)|_{[0,t]})$$

for every $\omega \in \Omega$. Define $\phi \colon \Lambda_T \to [0,T] \times \widehat{\Omega}_T^p$ as $\phi(\mathbb{X}|_{[0,t]}) = (t, \widetilde{\mathbb{X}}|_{[0,T]})$ where $\widetilde{\mathbb{X}}|_{[0,T]}$ denotes the stopped process defined in Definition 3.1 . Note that ϕ is continuous, thus measurable. Define $f \colon [0,T] \times \widehat{\Omega}_T^p \to \mathbb{R}$ as $f(t, \widehat{\mathbb{X}}) \coloneqq \mathbb{1}_{A_t}(\widehat{\mathbb{X}}|_{[0,t]})$. For fixed t, $\widehat{\mathbb{X}} \to \widehat{\mathbb{X}}|_{[0,t]}$ is continuous and $\widehat{\mathbb{X}}|_{[0,t]} \mapsto \mathbb{1}_{A_t}(\widehat{\mathbb{X}}|_{[0,t]})$ is measurable, therefore $\widehat{\mathbb{X}} \mapsto f(t, \widehat{\mathbb{X}})$ is measurable. For $n \in \mathbb{N}$, define $I_k^n \coloneqq [k/2^n T, (k+1)/2^n T)$ for $k = 0, \ldots, 2^n - 2$, $I_{2^n-1}^n \coloneqq [(2^n - 1)/2^n T, T]$ and $t_k^n \coloneqq k/2^n T$. Set

$$f_n(t, \widehat{\mathbb{X}}) \coloneqq \sum_{k=0}^{2^n-1} f(t_k^n, \widehat{\mathbb{X}}) \mathbb{1}_{I_k^n}(t)$$

which is measurable for every $n \in \mathbb{N}$. Set

$$\widetilde{f}(t, \widehat{\mathbb{X}}) \coloneqq \limsup_{m\to\infty} \limsup_{n\to\infty} f_n(t + 1/m, \widehat{\mathbb{X}}) \quad \text{and} \quad \theta(\widehat{\mathbb{X}}|_{[0,t]}) \coloneqq (\widetilde{f} \circ \phi)(\widehat{\mathbb{X}}|_{[0,t]}).$$

The map θ is thus measurable and satisfies

$$\theta(\widehat{\mathbb{X}}(\omega)|_{[0,t]}) = \limsup_{m\to\infty} \limsup_{n\to\infty} \sum_{k=0}^{2^n-1} \mathbb{1}_{\{\tau(\omega) \le t_k^n\}} \mathbb{1}_{I_k^n}(t + 1/m) = \mathbb{1}_{\{\tau(\omega) \le t\}}. \qquad \square$$

4. RANDOMIZED STOPPING TIMES

In Lemma 3.2 we have seen that a stopping time can be represented by a stopping policy of the form $\theta : \Lambda_T \to \{0, 1\}$. In the next step, we reverse the order and define stopping times associated to *continuous* stopping policies of the form $\theta : \Lambda_T \to \mathbb{R}$. Of course, for a given stopping time, there is no reason why it should be representable by a continuous stopping policy. Indeed, relevant stopping times – such as hitting times of even nice sets – are often discontinuous functions of the underlying path. We will see, however, that stopping times, in particular the optimal stopping times for our problem, can be approximated by stopping times induced by continuous policies, in the sense that the corresponding value functions converge. Later, in Section 5 and Section 7 we will see that this is also true when we restrict to the subclasses of continuous stopping policies given by linear functionals respectively deep neural networks applied to the signature.

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. In this and the following sections, $\widehat{\mathbb{X}}$ denotes a stochastic process in $\widehat{\Omega}_T^p$ and $Y \colon [0,T] \times \Omega \to \mathbb{R}$ is a real-valued continuous stochastic process adapted to the filtration $(\mathcal{F}_t)$, $\mathcal{F}_t = \sigma(\widehat{\mathbb{X}}_{0,s} : 0 \le s \le t)$. Denote by $\mathcal{S}$ the space of all $(\mathcal{F}_t)$-stopping times. We are trying to solve the optimal stopping problem for Y, i.e., in a financial context Y corresponds to a cash-flow process. For simplicity, we assume that $X_0 = 0$.

We consider *randomized stopping times*, which relax proper stopping times and lead to much more regular approximation problems. We note that similar techniques have been used in [BTW20] in the context of numerical methods for American option pricing.

Definition 4.1. We set $\mathcal{T} \coloneqq C(\Lambda_T, \mathbb{R})$ and call it the space of *continuous stopping policies*. Let Z be a non-negative random variable independent of $\widehat{\mathbb{X}}$ and such that $\mathbb{P}(Z = 0) = 0$. For a continuous stopping policy $\theta \in \mathcal{T}$, we define the *randomized stopping time* by

$$(4.1) \qquad \tau_\theta^r \coloneqq \inf\left\{ t \ge 0 \ : \ \int_0^{t\wedge T} \theta(\widehat{\mathbb{X}}|_{[0,s]})^2 \, ds \ge Z \right\}$$

where $\inf \emptyset = +\infty$.

Next we prove that stopping times can be approximated by randomized stopping times based on continuous stopping policies.

Proposition 4.2. *For every stopping time $\tau \in \mathcal{S}$, there exists a sequence $\theta_n \in \mathcal{T}$ such that the randomized stopping times $\tau_{\theta_n}^r$ satisfy $\tau_{\theta_n}^r \to \tau$ almost surely as $n \to \infty$. In particular, if $\mathbb{E}[\|Y\|_\infty] < \infty$, then*

$$\sup_{\theta \in \mathcal{T}} \mathbb{E}[Y_{\tau_\theta^r \wedge T}] = \sup_{\tau \in \mathcal{S}} \mathbb{E}[Y_{\tau \wedge T}].$$

Proof. Let τ be a stopping time. From Lemma 3.2, we know that there is a measurable map $\theta \colon \Lambda_T \to \{0,1\}$ such that

$$\theta(\widehat{\mathbb{X}}|_{[0,t]}) = \mathbb{1}_{\{\tau \le t\}}.$$

Using [Wiś94, Theorem 1], we can find a sequence of continuous functions $\widetilde{\theta}_n \in \mathcal{T}$ such that $\widetilde{\theta}_n(\widehat{\mathbb{X}}|_{[0,t]}) \to \mathbb{1}_{\{\tau \le t\}}$ almost surely w.r.t. $\lambda|_{[0,T]} \otimes \mathbb{P}$ where $\lambda|_{[0,T]}$ denotes the Lebesgue measure on $[0,T]$. W.l.o.g, we may assume that $0 \le \widetilde{\theta}_n \le 1$. Set $\theta_n := (2\widetilde{\theta}_n)^n$. Then

$$\lim_{n \to \infty} \theta_n(\widehat{\mathbb{X}}|_{[0,t]}) \to \begin{cases} +\infty & \text{if } t \ge \tau \\ 0 & \text{if } t < \tau. \end{cases}$$

It follows that $\tau_{\theta_n}^r \to \tau$ almost surely as $n \to \infty$. Using the dominated convergence theorem, this implies that

$$\sup_{\theta \in \mathcal{T}} \mathbb{E}[Y_{\tau_\theta^r \wedge T}] \ge \sup_{\tau \in \mathcal{S}} \mathbb{E}[Y_{\tau \wedge T}].$$

To show the converse inequality, take $\theta \in \mathcal{T}$. From independence,

$$\mathbb{E}[Y_{\tau_\theta^r \wedge T} \mid \widehat{\mathbb{X}}] = \int_0^\infty Y_{\tau_z \wedge T} \, \mathbb{P}_Z(dz)$$

where

$$\tau_z := \inf \left\{ t \ge 0 \, : \, \int_0^{t \wedge T} \theta(\widehat{\mathbb{X}}|_{[0,s]})^2 \, ds \ge z \right\}.$$

Note that this is a stopping time for every $z \ge 0$. Taking expectation, it follows that

$$\mathbb{E}[Y_{\tau_\theta^r \wedge T}] = \int_0^\infty \mathbb{E}[Y_{\tau_z \wedge T}] \, \mathbb{P}_Z(dz) \le \sup_{\tau \in \mathcal{S}} \mathbb{E}[Y_{\tau \wedge T}]$$

which implies the claim. $\qquad\qquad\square$

Note that we cannot generally assume that $\theta_n \to \theta$ implies $\tau_{\theta_n}^r \to \tau_\theta^r$, as is shown by the following counter-example. So even randomized stopping times are not continuous w.r.t. the underlying stopping policies.

Example 4.3. Consider $\vartheta, \vartheta_n \colon [0,3] \to [0,\infty)$ defined by

$$\vartheta(t) = \begin{cases} 1 - t & \text{if } t \in [0,1] \\ 0 & \text{if } t \in [1,2] \\ t - 2 & \text{if } t \in [2,3] \end{cases} \quad \text{and} \quad \vartheta_n(t) = \begin{cases} (1 - \frac{1}{n})(1 - t) & \text{if } t \in [0,1] \\ 0 & \text{if } t \in [1,2] \\ t - 2 & \text{if } t \in [2,3]. \end{cases}$$

Although $\vartheta_n \to \vartheta$ as $n \to \infty$, we have

$$\inf\left\{t \geq 0 \,:\, \int_0^{t \wedge 3} \vartheta(s)\, ds \geq \frac{1}{2}\right\} = 1 \quad \text{and} \quad \inf\left\{t \geq 0 \,:\, \int_0^{t \wedge 3} \vartheta_n(s)\, ds \geq \frac{1}{2}\right\} > 2$$

for all $n \geq 1$.

As mentioned above, randomized stopping times *regularize* the optimal stopping problem. Indeed, given a randomized stopping time τ_θ^r defined in terms of an independent random variable Z as in Definition 4.1, if we integrate the stopped process $Y_{\tau_\theta^r \wedge T}$ w.r.t. Z, we obtain a smooth function of θ – which is clearly not true without regularization, see Remark 4.5 below.

Proposition 4.4. *Let S be an $(\mathcal{F}_t)$-stopping time and let F_Z denote the cumulative distribution function of Z. Then*

$$\mathbb{E}[Y_{\tau_\theta^r \wedge S} \,|\, \widehat{\mathbb{X}}] = \int_0^S Y_t \, d\widetilde{F}(t) + Y_S(1 - \widetilde{F}(S)) = \int_0^S (1 - \widetilde{F}(t))\, dY_t + Y_0$$

where the second integral is implicitly defined by integration by parts and

$$\widetilde{F}(t) := F_Z\left(\int_0^t \theta(\widehat{\mathbb{X}}|_{[0,s]})^2 \, ds\right).$$

In particular, if Z has a density ϱ,

$$\mathbb{E}[Y_{\tau_\theta^r \wedge S}] = \mathbb{E}\left[\int_0^S Y_t \theta(\widehat{\mathbb{X}}|_{[0,t]})^2 \varrho\left(\int_0^t \theta(\widehat{\mathbb{X}}|_{[0,s]})^2 \, ds\right) dt + Y_S(1 - \widetilde{F}(S))\right].$$

Proof. Recall that $\tau_\theta^r \in [0, T] \cup \{\infty\}$. For $t \in [0, \infty)$, we have

$$\mathbb{P}(\tau_\theta^r \leq t \,|\, \widehat{\mathbb{X}}) = \mathbb{P}\left(\int_0^{t \wedge T} \theta(\widehat{\mathbb{X}}|_{[0,s]})^2 \, ds \geq Z \,|\, \widehat{\mathbb{X}}\right) = F_Z\left(\int_0^{t \wedge T} \theta(\widehat{\mathbb{X}}|_{[0,s]})^2 \, ds\right) = \widetilde{F}(t)$$

and

$$\mathbb{P}(\tau_\theta^r = \infty \,|\, \widehat{\mathbb{X}}) = \mathbb{P}\left(\int_0^T \theta(\widehat{\mathbb{X}}|_{[0,s]})^2 \, ds < Z \,|\, \widehat{\mathbb{X}}\right) = 1 - \widetilde{F}(T).$$

It follows that for $f \colon [0, \infty] \to \mathbb{R}$ integrable,

$$\mathbb{E}[f(\tau_\theta^r) \,|\, \widehat{\mathbb{X}}] = \int_0^T f(t) \, d\widetilde{F}(t) + f(\infty)(1 - \widetilde{F}(T))$$

and therefore

$$\mathbb{E}[Y_{\tau_\theta^r \wedge S} \,|\, \widehat{\mathbb{X}}] = \int_0^T Y_{t \wedge S} \, d\widetilde{F}(t) + Y_S(1 - \widetilde{F}(T)) = \int_0^S Y_t \, d\widetilde{F}(t) + Y_S(1 - \widetilde{F}(S)).$$

$\square$

Remark 4.5. In the deterministic case $Z = z > 0$ almost surely, we have

$$\widetilde{F}(t) = \mathbb{1}_{[z, \infty)}\left(\int_0^t \theta(\widehat{\mathbb{X}}|_{[0,s]})^2 \, ds\right),$$

thus

$$\mathbb{E}[Y_{\tau_\theta^r \wedge T}] = \mathbb{E}\left[\int_0^T \mathbb{1}_{[0,z)}\left(\int_0^t \theta(\widehat{\mathbb{X}}|_{[0,s]})^2\, ds\right)\, dY_t\right] + \mathbb{E}[Y_0].$$

5. Linear signature stopping policies

Using the regularization by randomization we will prove that it is enough to use stopping policies that are linear functionals of the signature.

Definition 5.1. The space of *linear signature stopping policies* $\mathcal{T}_{\text{sig}} \subset \mathcal{T}$ is defined as

$$\mathcal{T}_{\text{sig}} = \left\{\theta \in \mathcal{T} \,:\, \exists l \in T((\mathbb{R}^{1+d})^*) \text{ such that } \theta(\widehat{\mathbb{X}}|_{[0,t]}) = \langle l, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle \; \forall \widehat{\mathbb{X}}|_{[0,t]} \in \Lambda_T\right\}.$$

Note that every $l \in T((\mathbb{R}^{1+d})^*)$ defines a $\theta_l \in \mathcal{T}$ by setting $\theta_l(\widehat{\mathbb{X}}|_{[0,t]}) \coloneqq \langle l, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle$. Let Z be as in Definition 4.1, then we introduce the following notation for randomized stopping times associated to linear signature stopping policies

$$(5.1) \qquad \tau_l^r \coloneqq \tau_{\theta_l}^r = \inf\left\{t \geq 0 \,:\, \int_0^{t \wedge T} \langle l, \widehat{\mathbb{X}}_{0,s}^{<\infty}\rangle^2\, ds \geq Z\right\}.$$

As a consequence of the Stone-Weierstrass theorem, $\mathcal{T}_{\text{sig}}$ is dense in $\mathcal{T}$. More precisely, we have

Lemma 5.2. *Let $\mathbb{P}$ be a probability measure on $(\widehat{\Omega}_T^p, \mathcal{B}(\widehat{\Omega}_T^p))$. Then, for every $\varepsilon > 0$, there is a compact set $\mathcal{K} \subset \widehat{\Omega}_T^p$ such that*

(1) $\mathbb{P}(\mathcal{K}) > 1 - \varepsilon$,

(2) $\mathcal{T}_{\text{sig}}$, *restricted of $\mathcal{K}$, is dense in $\mathcal{T}$. More precisely, for every $\theta \in \mathcal{T}$ there is a sequence $\theta_n \in \mathcal{T}_{\text{sig}}$ such that*

$$\sup_{\widehat{\mathbb{X}} \in \mathcal{K};\; t \in [0,T]} |\theta_n(\widehat{\mathbb{X}}|_{[0,t]}) - \theta(\widehat{\mathbb{X}}|_{[0,t]})| \to 0$$

as $n \to \infty$.

Proof. [KLPA20, Lemma B.3]. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

The main result of this section is the following

Proposition 5.3. *Assume that Z has a continuous density ϱ and that $\mathbb{E}[\|Y\|_\infty] < \infty$. Then*

$$\sup_{\theta \in \mathcal{T}} \mathbb{E}[Y_{\tau_\theta^r \wedge T}] = \sup_{\theta \in \mathcal{T}_{\text{sig}}} \mathbb{E}[Y_{\tau_\theta^r \wedge T}].$$

It follows that

$$\sup_{\theta \in \mathcal{T}} \mathbb{E}[Y_{\tau_\theta^r \wedge T}] = \sup_{l \in T((\mathbb{R}^{1+d})^*)} \mathbb{E}[Y_{\tau_l^r \wedge T}].$$

Proof. It is enough to show that $\sup_{\theta \in \mathcal{T}} \mathbb{E}[Y_{\tau_\theta^r \wedge T}] \leq \sup_{\theta \in \mathcal{T}_{\mathrm{sig}}} \mathbb{E}[Y_{\tau_\theta^r \wedge T}]$. Let $\theta \in \mathcal{T}$. From Lemma 5.2, we know that for every $\varepsilon > 0$, there is a compact set $\mathcal{K} \subset \widehat{\Omega}_T^p$ such that for $A := \{\widehat{\mathbb{X}} \in \mathcal{K}\}$, we have $\mathbb{P}(A) \geq 1 - \varepsilon$, and a sequence $\theta_n \in \mathcal{T}_{\mathrm{sig}}$ such that

$$(5.2) \qquad \lim_{n \to \infty} \sup_{\mathbb{X} \in \mathcal{K}; t \in [0,T]} |\theta_n(\mathbb{X}|_{[0,t]}) - \theta(\mathbb{X}|_{[0,t]})| = 0.$$

Let $\mathcal{K}$ be such a compact set, the precise choice will be made later. Set

$$\widetilde{F}_n(t) = F_Z\left(\int_0^t \theta_n(\widehat{\mathbb{X}}|_{[0,s]})^2\, ds\right) \quad \text{and} \quad \widetilde{F}(t) = F_Z\left(\int_0^t \theta(\widehat{\mathbb{X}}|_{[0,s]})^2\, ds\right).$$

Then,

$$|\mathbb{E}[Y_T(1 - \widetilde{F}_n(T))\,;\, A] - \mathbb{E}[Y_T(1 - \widetilde{F}(T))\,;\, A]|$$
$$\leq \mathbb{E}[|Y_T||\widetilde{F}_n(T) - \widetilde{F}(T)|\,;\, A]$$
$$\leq \mathbb{E}[|Y_T|] \sup_{\mathbb{X} \in \mathcal{K}} \left| F_Z\left(\int_0^T \theta_n(\mathbb{X}|_{[0,s]})^2\, ds\right) - F_Z\left(\int_0^T \theta(\mathbb{X}|_{[0,s]})^2\, ds\right)\right|.$$

Since F_Z is continuous and uniformly continuous on compact sets,

$$\sup_{\mathbb{X} \in \mathcal{K}} \left| F_Z\left(\int_0^T \theta_n(\mathbb{X}|_{[0,s]})^2\, ds\right) - F_Z\left(\int_0^T \theta(\mathbb{X}|_{[0,s]})^2\, ds\right)\right| \to 0$$

as $n \to \infty$. Indeed: we first show that

$$(5.3) \qquad \sup_{\mathbb{X} \in \mathcal{K}; t \in [0,T]} |\theta_n(\mathbb{X}|_{[0,t]})^2 - \theta(\mathbb{X}|_{[0,t]})^2| \to 0$$

as $n \to \infty$. Since $\sup_{n \geq 1} \sup_{\mathbb{X} \in \mathcal{K}; t \in [0,T]} |\theta_n(\mathbb{X}|_{[0,t]})| < \infty$, the functions θ and θ_n take their values in a compact set, hence (5.3) follows from (5.2). Property (5.2) also implies that

$$\sup_{\mathbb{X} \in \mathcal{K}} \left| \int_0^T \theta_n(\mathbb{X}|_{[0,s]})^2\, ds - \int_0^T \theta(\mathbb{X}|_{[0,s]})^2\, ds\right| \to 0$$

as $n \to \infty$. Using continuity of F_Z and uniform continuity on compact sets implies the claim. It follows that

$$\lim_{n \to \infty} |\mathbb{E}[Y_T(1 - \widetilde{F}_n(T))\,;\, A] - \mathbb{E}[Y_T(1 - \widetilde{F}(T))\,;\, A]| = 0.$$

Since $|\widetilde{F}_n(T) - \widetilde{F}(T)| \leq 2$,

$$|\mathbb{E}[Y_T(1 - \widetilde{F}_n(T))\,;\, A^c] - \mathbb{E}[Y_T(1 - \widetilde{F}(T))\,;\, A^c]| \leq 2\mathbb{E}[|Y_T|\,;\, A^c]$$

and this quantity can me made arbitrarily small by the choice of $\mathcal{K}$.

With the same arguments, we can show that

$$\left| \mathbb{E}\left[\int_0^T Y_t \left(\theta_n(\widehat{\mathbb{X}}|_{[0,t]})^2 \varrho\left(\int_0^t \theta_n(\widehat{\mathbb{X}}|_{[0,s]})^2\, ds\right) - \theta(\widehat{\mathbb{X}}|_{[0,t]})^2 \varrho\left(\int_0^t \theta(\widehat{\mathbb{X}}|_{[0,s]})^2\, ds\right)\right) dt\right]\right|$$
$$\to 0$$

as $n \to \infty$ which implies the claim. $\qquad\square$

Finally, we note that we do not need randomization for the approximation by stopping times based on signature stopping policies to work.

Definition 5.4. For $l \in T((\mathbb{R}^{1+d})^*)$ define the hitting time of the signature against the half-plane orthogonal to l by

$$\tau_l := \inf \left\{ t \in [0, T] \ : \ \langle l, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \geq 1 \right\}.$$

Proposition 5.5. *Given* $\mathbb{E}[\|Y\|_\infty] < \infty$, *we have*

$$\sup_{l \in T((\mathbb{R}^{1+d})^*)} \mathbb{E}[Y_{\tau_l \wedge T}] = \sup_{\tau \in \mathcal{S}} \mathbb{E}[Y_{\tau \wedge T}].$$

Proof. Using Proposition 4.2 and 5.3, it suffices to show that

$$\sup_{l \in T((\mathbb{R}^{1+d})^*)} \mathbb{E}[Y_{\tau_l^r \wedge T}] \leq \sup_{l \in T((\mathbb{R}^{1+d})^*)} \mathbb{E}[Y_{\tau_l \wedge T}].$$

Choose $l \in T((\mathbb{R}^{1+d})^*)$. Then

$$\mathbb{E}[Y_{\tau_l^r \wedge T} \mid \widehat{\mathbb{X}}] = \int_0^\infty Y_{\tau_z \wedge T} \, \mathbb{P}_Z(dz)$$

where

$$\tau_z := \inf \left\{ t \geq 0 \ : \ \int_0^{t \wedge T} \langle l, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle^2 \, ds \geq z \right\} = \inf \left\{ t \in [0, T] \ : \ \langle (l \shuffle l) 1/z, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \geq 1 \right\}$$

which is a signature hitting time for every $z > 0$ in the sense of the above definition. Taking expectation, we obtain

$$\mathbb{E}[Y_{\tau_l^r \wedge T}] = \int_0^\infty \mathbb{E}[Y_{\tau_z \wedge T}] \, \mathbb{P}_Z(dz) \leq \sup_{\ell \in T((\mathbb{R}^d)^*)} \mathbb{E}[Y_{\tau_\ell \wedge T}]$$

as claimed. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ □

Remark 5.6. In the case of X being a standard Markov process in $\mathbb{R}^d$ and $Y_t = G(t, X_t)$ for a continuous function G, it is known that

$$\sup_{\tau \in \mathcal{S}} \mathbb{E}[Y_{\tau \wedge T}] = \sup_{\tau \in \mathfrak{D}} \mathbb{E}[Y_{\tau \wedge T}]$$

where $\mathfrak{D}$ denotes the set of all hitting times of closed sets in $\mathbb{R}^{1+d}$ of the process $t \mapsto (t, X_t)$ [Shi08, Corollary 3 on p. 129]. Our Theorem can be seen as an extension of this classical result to non-Markovian processes.

6. LINEARIZATION OF THE OPTIMAL STOPPING PROBLEM

In this section, we will study randomized signature stopping times associated to linear stopping policies in $\mathcal{T}_{\text{sig}}$ and a exponentially distributed random variable $Z \sim \text{Exp}(1)$. From Proposition 4.4 we have
(6.1)

$$\sup_{\tau \in \mathcal{S}} \mathbb{E}[Y_{\tau \wedge T}] = \sup_{l \in T((\mathbb{R}^{1+d})^*)} \mathbb{E}[Y_{\tau_l^r \wedge T}] = \sup_{l \in T((\mathbb{R}^{1+d})^*)} \mathbb{E}\left[\int_0^T \exp\left(-\int_0^t \langle l, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle^2 \, ds \right) dY_t \right] + \mathbb{E}[Y_0].$$

Recall that for group-like elements $\mathbf{a} \in G(V)$ as defined in (2.2), polynomials of linear functionals in $\mathbf{a}$ can be expressed in terms of shuffle products of the linear functionals themselves, see (2.3). Consider the main term on the right-hand side of (6.1):

- The innermost term $\langle l, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle^2$ is a polynomial of a linear functional of the signature. It can, therefore, be expressed as a linear functional of the signature, more precisely, $\langle l, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle^2 = \langle l \sqcup\!\sqcup l, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle$.

- Given an element of the signature, integrating against a component of the underlying path produces another element of the signature. Concretely,

$$\int_0^t \langle l, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle^2 ds = \langle (l \sqcup\!\sqcup l)\mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle,$$

  recalling that the time-component of our driving path $\widehat{X}_t = (t, X_t)$ was associated with the letter $\mathbf{1}$, see Remark 2.2.

- Next we need to apply the exponential function to $\langle (l \sqcup\!\sqcup l)\mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$. Unfortunately, the exponential function is not a polynomial, so we cannot directly apply the shuffle product. However, as we shall see below, there is a corresponding *exponential shuffle*, which comes with certain restrictions. Nonetheless, we shall see that we will still obtain a linear functional of the signature for our purposes.

- Finally, we integrate against Y and take the expectation. **If** Y can itself be represented as a linear functional of the signature, integrating another linear functional of the signature against Y will result in yet another linear functional of the signature. In this case, we can finally interchange the expectation, and the right-hand side of (6.1) can be represented as a sup over a linear functional of the *expected signature* $\mathbb{E}[\widehat{\mathbb{X}}_{0,T}^{<\infty}]$ of $\widehat{X}$.

In the remainder of this section, we will follow through with this program.

We start with a definition of an exponential function based on the shuffle product.

Definition 6.1. For $l \in T(V^*)$ with $l = a_0 \varnothing + \widetilde{l}$ and $\langle \widetilde{l}, \mathbf{1} \rangle = 0$ we define the *exponential shuffle*

$$(6.2) \qquad \exp^{\sqcup\!\sqcup}(l) := \exp(a_0) \exp^{\sqcup\!\sqcup}(\widetilde{l}), \quad \text{where} \quad \exp^{\sqcup\!\sqcup}(\widetilde{l}) := \sum_{r=0}^{\infty} \frac{1}{r!} \widetilde{l}^{\sqcup\!\sqcup r}.$$

Since obviously $\pi_{\leq N}(\widetilde{l}^{\sqcup\!\sqcup r}) = 0$ for $r > N$, the infinite sum is well defined as an element in the extended tensor algebra $T((V^*))$. One may easily check that $\exp^{\sqcup\!\sqcup}(\widetilde{l}_1 + \widetilde{l}_2) = \exp^{\sqcup\!\sqcup}(\widetilde{l}_1) \exp^{\sqcup\!\sqcup}(\widetilde{l}_2)$ for $\widetilde{l}_1, \widetilde{l}_2 \in T(V^*)$ such that $\langle \widetilde{l}_1, \mathbf{1} \rangle = \langle \widetilde{l}_2, \mathbf{1} \rangle = 0$. Thus, in particular, one has

$$(6.3) \qquad \exp^{\sqcup\!\sqcup}(l_1 + l_2) = \exp^{\sqcup\!\sqcup}(l_1) \exp^{\sqcup\!\sqcup}(l_2) \quad \text{for all} \ l_1, l_2 \in T(V^*).$$

Note also that

$$\exp^{\sqcup\!\sqcup}(l) = \sum_{r=0}^{\infty} \frac{1}{r!} l^{\sqcup\!\sqcup r}$$

for every $l \in T(V^*)$.

We can now prove that the exponential shuffle linearizes the exponential function for group-like elements. In this context, keep in mind that $\exp^{\sqcup\!\sqcup}(l) \notin T(V^*)$ and, hence, it is not a well-defined linear functional on $T((V))$. It is, however, trivially well-defined as a linear functional on $T(V)$, and, hence, can be applied to any projection $\pi_{\leq N}(\mathbf{v})$, $\mathbf{v} \in T((V))$. In addition, as the lemma shows, we can apply $\exp^{\sqcup\!\sqcup}(l)$ to group-like elements.

Lemma 6.2. *Let $l \in T(V^*)$ and $\mathbf{g} \in G(V)$. One then has*

$$|\exp(\langle l, \mathbf{g} \rangle) - \langle \exp^{\sqcup\!\sqcup}(l), \pi_{\leq N}(\mathbf{g}) \rangle| \leq 4 \exp(\langle l, 1 \rangle) \frac{\big(|l| \|\pi_{\leq \deg(l)}(\mathbf{g})\|\big)^{\lfloor N/\deg(l) \rfloor + 1}}{(\lfloor N/\deg(l) \rfloor + 1)!}$$

for $N > 2\deg(l)|l| \|\pi_{\leq \deg(l)}(\mathbf{g})\|$.

Proof. Let us write $l = a_0 \varnothing + \widetilde{l}$ with $\langle \widetilde{l}, \mathbf{1} \rangle = 0$, where for mutually different words $w_1, ..., w_n$,

$$\widetilde{l} = \lambda_1 w_1 + ... + \lambda_n w_n, \quad \text{and set} \quad M := \deg(l) = \max_{1 \leq i \leq n} \deg(w_i), \quad m := \min_{1 \leq i \leq n} \deg(w_i) \geq 1.$$

We then have

$$\pi_{\leq N}\left(\exp^{\sqcup\!\sqcup}(\widetilde{l})\right) = \sum_{r=0}^{\lfloor N/M \rfloor} \frac{\widetilde{l}^{\sqcup\!\sqcup r}}{r!} + \pi_N \left(\sum_{r=\lfloor N/\deg(l) \rfloor + 1}^{\lfloor N/m \rfloor} \frac{\widetilde{l}^{\sqcup\!\sqcup r}}{r!} \right) \in T^N(V^*) \quad \text{for any} \quad N \geq 1.$$

Hence,

$$\langle \pi_{\leq N}\left(\exp^{\sqcup\!\sqcup}(l)\right), \mathbf{g} \rangle = \exp(a_0) \sum_{r=0}^{\lfloor N/M \rfloor} \frac{\langle \widetilde{l}, \mathbf{g} \rangle^r}{r!} + \exp(a_0) \sum_{r=\lfloor N/M \rfloor + 1}^{\lfloor N/m \rfloor} \frac{1}{r!} \langle \pi_N\left(\widetilde{l}^{\sqcup\!\sqcup r}\right), \mathbf{g} \rangle$$

$$=: \exp(a_0) \sum_{r=0}^{\lfloor N/M \rfloor} \frac{\langle \widetilde{l}, \mathbf{g} \rangle^r}{r!} + R_N^{(1)},$$

and since $\mathbf{g} \in G(V)$ it holds that

$$\langle \pi_{\leq N}\left(\widetilde{l}^{\sqcup\!\sqcup r}\right), \mathbf{g} \rangle = \sum_{\substack{i_1, ..., i_r = 1 \\ \deg(w_{i_1}) + ... + \deg(w_{i_r}) \leq N}}^{n} \lambda_{i_1} ... \lambda_{i_r} \langle w_{i_1} \sqcup\!\sqcup ... \sqcup\!\sqcup w_{i_r}, \mathbf{g} \rangle$$

$$= \sum_{\substack{i_1, ..., i_r = 1 \\ \deg(w_{i_1}) + ... + \deg(w_{i_r}) \leq N}}^{n} \langle \lambda_{i_1} w_{i_1}, \mathbf{g} \rangle ... \langle \lambda_{i_r} w_{i_r}, \mathbf{g} \rangle .$$

Hence we have

$$\left| \langle \pi_{\leq N}\left(\widetilde{l}^{\sqcup\!\sqcup r}\right), \mathbf{g} \rangle \right| \leq \sum_{i_1, ..., i_r = 1}^{n} |\langle \lambda_{i_1} w_{i_1}, \mathbf{g} \rangle| ... |\langle \lambda_{i_r} w_{i_r}, \mathbf{g} \rangle| = \left(\sum_{i=1}^{n} |\langle \lambda_i w_i, \mathbf{g} \rangle| \right)^r$$

$$\leq |\widetilde{l}|^r \left| \pi_{\leq \deg(l)}(\mathbf{g}) \right|^r, \quad \text{and so}$$

$$\left| R_N^{(1)} \right| \leq \exp(a_0) \sum_{r=\lfloor N/M \rfloor + 1}^{\infty} \frac{|\widetilde{l}|^r \left| \pi_{\leq \deg(l)}(\mathbf{g}) \right|^r}{r!} \leq 2 \exp(a_0) \frac{\big(|\widetilde{l}| \|\pi_{\leq \deg(l)}(\mathbf{g})\|\big)^{\lfloor N/M \rfloor + 1}}{(\lfloor N/M \rfloor + 1)!}$$

for $N > N_{l,\mathbf{g}} := 2M||\widetilde{l}|\pi_{\leq \deg(l)}(\mathbf{g})|$. One further has (note that $\mathbf{g}_0 = \mathbf{1}$), due to a similar estimation,

$$\exp(\langle l, \mathbf{g} \rangle) = \exp(a_0) \exp(\langle \widetilde{l}, \mathbf{g} \rangle) = \exp(a_0) \sum_{r=0}^{\lfloor N/M \rfloor} \frac{\langle \widetilde{l}, \mathbf{g} \rangle^r}{r!} + R_N^{(2)} \quad \text{with}$$

$$\left| R_N^{(2)} \right| \leq 2 \exp(a_0) \frac{\left| \langle \widetilde{l}, \mathbf{g} \rangle \right|^{\lfloor N/M \rfloor + 1}}{(\lfloor N/M \rfloor + 1)!} \leq 2 \exp(a_0) \frac{\left(\left| \widetilde{l} \right| |\pi_{\leq \deg(l)}(\mathbf{g})| \right)^{\lfloor N/M \rfloor + 1}}{(\lfloor N/M \rfloor + 1)!}$$

for $N > N_{l,\mathbf{g}}$. Finally, by noting that $\langle \pi_{\leq N} (\exp^{\sqcup\!\sqcup}(l)), \mathbf{g} \rangle = \langle \exp^{\sqcup\!\sqcup}(l), \pi_{\leq N}(\mathbf{g}) \rangle$, and then taking all together we obtain the stated result. $\qquad\square$

Remark 6.3. The equation $\langle \exp^{\sqcup\!\sqcup}(l), \mathbf{g} \rangle = \exp(\langle l, \mathbf{g} \rangle)$ is confusing at first glance, because $\mathbf{g} \mapsto \langle \exp^{\sqcup\!\sqcup}(l), \mathbf{g} \rangle$ *seems* linear, whereas $\mathbf{g} \mapsto \exp(\langle l, \mathbf{g} \rangle)$ clearly is not. Note, however, that $\exp^{\sqcup\!\sqcup}(l) \in T((V^*))$ and, hence, does not define a linear map on $T((V))$. Indeed, the group $G(V)$ is not closed under linear combination, and, hence, Lemma 6.2 does simply not apply to a linear combination of elements $\mathbf{g}_1, \mathbf{g}_2 \in G(V)$.

The exponential shuffle satisfies a differential equation, which we shall use later. Note that terms of the form $\langle \exp^{\sqcup\!\sqcup}(l1), \widehat{\mathbb{X}}_{0,t}^{\leq N} \rangle$ are (classically) differentiable in t.

Lemma 6.4. *For every polynomial $l = \lambda_1 w_1 + \ldots + \lambda_n w_n \in T((\mathbb{R}^{1+d})^*)$,*

$$\frac{d}{dt} \langle \exp^{\sqcup\!\sqcup}(l1), \widehat{\mathbb{X}}_{0,t}^{\leq N} \rangle = \sum_{i=1}^{n} \langle \lambda_i w_i, \widehat{\mathbb{X}}_{0,t}^{< \infty} \rangle \langle \exp^{\sqcup\!\sqcup}(l1), \widehat{\mathbb{X}}_{0,t}^{\leq N - \deg(w_i) - 1} \rangle.$$

Proof. Note that

$$\frac{d}{dt} \langle w1, \widehat{\mathbb{X}}_{0,t}^{< \infty} \rangle = \frac{d}{dt} \int_0^t \langle w, \widehat{\mathbb{X}}_{0,s}^{< \infty} \rangle \, ds = \langle w, \widehat{\mathbb{X}}_{0,t}^{< \infty} \rangle$$

for every word w. Hence, for $l = \lambda_1 w_1 + \ldots + \lambda_n w_n$, one always has $\langle l1, \mathbf{1} \rangle = 0$ and so by (6.2),

$$\frac{d}{dt} \langle \exp^{\sqcup\!\sqcup}(l1), \widehat{\mathbb{X}}_{0,t}^{\leq N} \rangle$$

$$= \frac{d}{dt} \sum_{0 \leq k_1 \deg(w_1 1) + \ldots + k_n \deg(w_n 1) \leq N} \frac{\langle \lambda_1 w_1 1, \widehat{\mathbb{X}}_{0,t}^{< \infty} \rangle^{k_1}}{k_1!} \cdots \frac{\langle \lambda_n w_n 1, \widehat{\mathbb{X}}_{0,t}^{< \infty} \rangle^{k_n}}{k_n!}$$

$$= \sum_{0 \leq k_1 \deg(w_1 1) + \ldots + k_n \deg(w_n 1) \leq N} \langle \lambda_1 w_1, \widehat{\mathbb{X}}_{0,t}^{< \infty} \rangle \frac{\langle \lambda_1 w_1 1, \widehat{\mathbb{X}}_{0,t}^{< \infty} \rangle^{k_1 - 1}}{(k_1 - 1)!} \frac{\langle \lambda_2 w_2 1, \widehat{\mathbb{X}}_{0,t}^{< \infty} \rangle^{k_2}}{k_2!} \cdots \frac{\langle \lambda_n w_n 1, \widehat{\mathbb{X}}_{0,t}^{< \infty} \rangle^{k_n}}{k_n!}$$

$$+ \ldots + \langle \lambda_n w_n, \widehat{\mathbb{X}}_{0,t}^{< \infty} \rangle \frac{\langle \lambda_1 w_1 1, \widehat{\mathbb{X}}_{0,t}^{< \infty} \rangle^{k_1}}{k_1!} \cdots \frac{\langle \lambda_{n-1} w_{n-1} 1, \widehat{\mathbb{X}}_{0,t}^{< \infty} \rangle^{k_{n-1}}}{k_{n-1}!} \frac{\langle \lambda_n w_n 1, \widehat{\mathbb{X}}_{0,t}^{< \infty} \rangle^{k_n - 1}}{(k_n - 1)!}$$

and

$$\sum_{0 \le k_1 \deg(w_1 1)+\ldots+k_n \deg(w_n 1) \le N} \frac{\langle \lambda_1 w_1 1, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle^{k_1-1}}{(k_1-1)!} \frac{\langle \lambda_2 w_2 1, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle^{k_2}}{k_2!} \cdots \frac{\langle \lambda_n w_n 1, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle^{k_n}}{k_n!}$$

$$= \sum_{0 \le (k_1+1) \deg(w_1 1)+\ldots+k_n \deg(w_n 1) \le N} \frac{\langle \lambda_1 w_1 1, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle^{k_1}}{k_1!} \frac{\langle \lambda_2 w_2 1, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle^{k_2}}{k_2!} \cdots \frac{\langle \lambda_n w_n 1, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle^{k_n}}{k_n!}$$

$$= \langle \exp^{\sqcup\!\sqcup}(l1), \widehat{\mathbb{X}}_{0,t}^{\le N - \deg(w_1 1)} \rangle. \qquad \square$$

We are now ready to formulate the main result of this section. Consider the optimization problem (6.1), which we modify by expressing the exponential by the exponential shuffle. Then we obtain convergence to the value of the optimal stopping problem. The proof requires us to localize w.r.t. the rough path metric. Other than that, the below formulation is now essentially implementable: In particular, the result is formulated in terms of truncated signatures, which is necessary also from a numerical point of view.

Proposition 6.5. *For given $\kappa > 0$, we define the stopping time*

$$S = S_\kappa = \inf\{t \ge 0 \,:\, \|\widehat{\mathbb{X}}\|_{p-var;[0,t]} \ge \kappa\} \wedge T.$$

Assume $Z \sim \mathrm{Exp}(1)$ and $\mathbb{E}\left[\|Y\|_\infty\right] < \infty$. Then

(6.4)

$$\sup_{l \in T((\mathbb{R}^{1+d})^*)} \mathbb{E}[Y_{\tau_l^r \wedge T}] = \lim_{\kappa \to \infty} \lim_{K \to \infty} \lim_{N \to \infty} \sup_{|l|+\deg(l) \le K} \mathbb{E}\left[\int_0^{S_\kappa} \langle \exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)1), \widehat{\mathbb{X}}_{0,t}^{\le N} \rangle \, dY_t\right] + \mathbb{E}[Y_0]$$

where the first two limit signs may be interchanged.

Proof. To ease notation, assume that $Y_0 = 0$. Since

$$|Y_{\tau_l \wedge T} - Y_{\tau_l \wedge S}| \le \sup_{|t-s| \le T-S} |Y_t - Y_s| \to 0$$

for every l as $\kappa \to \infty$ and

$$\lim_{K \to \infty} \sup_{|l|+\deg(l) \le K} \mathbb{E}[Y_{\tau_l \wedge \widehat{S}}] = \sup_{l \in T((\mathbb{R}^{1+d})^*)} \mathbb{E}[Y_{\tau_l \wedge \widehat{S}}],$$

with $\widehat{S}$ being either S or T, it follows that

(6.5) $$\sup_{l \in T((\mathbb{R}^{1+d})^*)} \mathbb{E}[Y_{\tau_l \wedge T}] = \lim_{\kappa \to \infty} \lim_{K \to \infty} \sup_{|l|+\deg(l) \le K} \mathbb{E}[Y_{\tau_l \wedge S}] = \lim_{K \to \infty} \lim_{\kappa \to \infty} \sup_{|l|+\deg(l) \le K} \mathbb{E}[Y_{\tau_l \wedge S}].$$

Now fix κ, K and l with $|l| + \deg(l) \le K$. Recall the estimate

$$\left|\int_0^T f(s) \, dg(s)\right| \le T\|f'\|_\infty \|g\|_\infty + |f(T)g(T) - f(0)g(0)|.$$

Note that

$$\exp\left(-\int_0^t \langle l, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle^2 \, ds\right) = \exp(-\langle (l \sqcup\!\sqcup l)1, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle).$$

Fix N. Then

$$\left| \mathbb{E}\left[\int_0^S \exp(-\langle (l \sqcup\!\sqcup l)\mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle)\, dY_t \right] - \mathbb{E}\left[\int_0^S \langle \exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N} \rangle\, dY_t \right] \right|$$
$$\leq (1+T)\mathbb{E}\left[\|Y\|_\infty \| \exp(-\langle (l \sqcup\!\sqcup l)\mathbf{1}, \widehat{\mathbb{X}}_{0,\cdot}^{<\infty} \rangle) - \langle \exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1}), \widehat{\mathbb{X}}_{0,\cdot}^{\leq N} \rangle \|_{\mathcal{C}^1[0,S]} \right].$$

Using Lemma 6.2,

$$\| \exp(-\langle (l \sqcup\!\sqcup l)\mathbf{1}, \widehat{\mathbb{X}}_{0,\cdot}^{<\infty} \rangle) - \langle \exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1}), \widehat{\mathbb{X}}_{0,\cdot}^{\leq N} \rangle \|_{\infty;[0,S]} \leq 4 \sup_{t \in [0,S]} \frac{(|(l \sqcup\!\sqcup l)\mathbf{1}||\widehat{\mathbb{X}}_{0,t}^{\leq 2\deg(l)+1}|)^M}{M!}.$$

where $M = \lfloor N/(2\deg(l)+1) \rfloor + 1$ provided N is sufficiently large. Clearly, $|(l \sqcup\!\sqcup l)\mathbf{1}| \leq C_K$. Using Lyons' Extension theorem [LCL07, Theorem 3.7], we can estimate

$$\sup_{t \in [0,S]} |\widehat{\mathbb{X}}_{0,t}^{\leq 2\deg(l)+1}| \leq \|\widehat{\mathbb{X}}^{\leq 2\deg(l)+1}\|_{p-\mathrm{var};[0,S]} \leq C(1+\|\widehat{\mathbb{X}}\|_{p-\mathrm{var};[0,S]})^{2K+1} = C(1+\kappa)^{2K+1}.$$

Therefore, we obtain an estimate of the form

$$\| \exp(-\langle (l \sqcup\!\sqcup l)\mathbf{1}, \widehat{\mathbb{X}}_{0,\cdot}^{<\infty} \rangle) - \langle \exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1}), \widehat{\mathbb{X}}_{0,\cdot}^{\leq N} \rangle \|_{\infty;[0,S]} \leq \frac{C^M}{M!}$$

for a deterministic constant C.

Next, we consider the derivatives. Set $\widetilde{l} = -(l \sqcup\!\sqcup l)$ and assume $\widetilde{l} = \lambda_1 w_1 + \ldots + \lambda_k w_k$. Clearly,

$$\frac{d}{dt} \exp(\langle \widetilde{l}\mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle) = \langle \widetilde{l}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \exp(\langle \widetilde{l}\mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle)$$

and Lemma 6.4 shows that

$$\frac{d}{dt} \langle \exp^{\sqcup\!\sqcup}(\widetilde{l}\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N} \rangle = \sum_{i=1}^k \langle \lambda_i w_i, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \langle \exp^{\sqcup\!\sqcup}(\widetilde{l}\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N-\deg(w_i)-1} \rangle.$$

Thus for $t \in [0,S]$,

$$\left| \frac{d}{dt} \left(\exp(\langle \widetilde{l}\mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle) - \langle \exp^{\sqcup\!\sqcup}(\widetilde{l}\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N} \rangle \right) \right|$$
$$\leq \sum_{i=1}^k |\langle \lambda_i w_i, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle| \left| \exp(\langle \widetilde{l}\mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle) - \langle \exp^{\sqcup\!\sqcup}(\widetilde{l}\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N-\deg(w_i)-1} \rangle \right|.$$

Using Lyons' Extension theorem,

$$|\langle \lambda_i w_i, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle| \leq C|\lambda_i| \|\widehat{\mathbb{X}}\|_{p-\mathrm{var};[0,S]}^{\deg(w_i)} \leq C|\lambda_i|$$

for a deterministic constant $C > 0$. Lemma 6.2 implies that for N sufficiently large,

$$\left| \exp(\langle \widetilde{l}\mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle) - \langle \exp^{\sqcup\!\sqcup}(\widetilde{l}\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N-\deg(w_i)-1} \rangle \right| \leq \frac{C^M}{M!}$$

for a deterministic constant $C > 0$ and $M \to \infty$ as $N \to \infty$. It follows that also

$$\left\| \frac{d}{dt} \left(\exp(-\langle (l \sqcup\!\sqcup l)\mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle) - \langle \exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N} \rangle \right) \right\|_{\infty;[0,S]} \leq \frac{C^M}{M!}.$$

This implies that

$$\sup_{|l|+\deg(l)\leq K} \left| \mathbb{E}\left[\int_0^S \exp(-\langle (l \shuffle l)\mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle)\, dY_t \right] - \mathbb{E}\left[\int_0^S \langle \exp^{\shuffle}(-(l \shuffle l)\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N}\rangle\, dY_t \right] \right| \to 0$$

as $N \to \infty$ and, in particular,

$$\lim_{N\to\infty} \sup_{|l|+\deg(l)\leq K} \mathbb{E}\left[\int_0^S \langle \exp^{\shuffle}(-(l \shuffle l)\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N}\rangle\, dY_t \right]$$

$$= \sup_{|l|+\deg(l)\leq K} \mathbb{E}\left[\int_0^S \exp(-\langle (l \shuffle l)\mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle)\, dY_t \right] = \sup_{|l|+\deg(l)\leq K} \mathbb{E}[Y_{\tau_l^r \wedge S}].$$

Together with (6.5), this proves (6.4). $\square$

Often, one is interested to solve the stopping problem for specific functionals of the underlying process X. In the next corollary, we consider a particular example. To simplify the exposition, we will consider the case $d = 1$ only. The generalization to arbitrary dimensions d is straightforward.

Corollary 6.6. *Assume $d = 1$ and that*

$$Y_t = G(X_t) + \int_0^t L(X_s)\, ds$$

for polynomials G and L. Then

$$\sup_{l\in T((\mathbb{R}^{1+d})^*)} \mathbb{E}[Y_{\tau_l^r \wedge T}]$$

$$= \lim_{\kappa\to\infty} \lim_{K\to\infty} \lim_{N\to\infty} \sup_{|l|+\deg(l)\leq K} \langle (\exp^{\shuffle}(-(l \shuffle l)\mathbf{1}) \shuffle G^{'\shuffle}(2))2 + (\exp^{\shuffle}(-(l \shuffle l)\mathbf{1}) \shuffle L^{\shuffle}(2))1, \mathbb{E}[\widehat{\mathbb{X}}_{0,S}^{\leq N}]\rangle$$

$$+ \mathbb{E}[Y_0].$$

In particular, if $d = 1$ and $X_0 = 0$,

$$(6.6) \qquad \sup_{l\in T((\mathbb{R}^2)^*)} \mathbb{E}[X_{\tau_l^r \wedge T}] = \lim_{\kappa\to\infty} \lim_{K\to\infty} \lim_{N\to\infty} \sup_{|l|+\deg(l)\leq K} \langle \exp^{\shuffle}(-(l \shuffle l)\mathbf{1})2, \mathbb{E}[\widehat{\mathbb{X}}_{0,S}^{\leq N}]\rangle.$$

Proof. We have

$$\int_0^S \langle \exp^{\shuffle}(-(l \shuffle l)\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N}\rangle\, dY_t = \int_0^S \langle \exp^{\shuffle}(-(l \shuffle l)\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N}\rangle G'(X_t)\, dX_t$$

$$+ \int_0^S \langle \exp^{\shuffle}(-(l \shuffle l)\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N}\rangle L(X_t)\, dt.$$

Since G' is a polynomial,

$$
\begin{aligned}
\int_0^S \langle \exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N} \rangle G'(X_t) \, dX_t &= \int_0^S \langle \exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N} \rangle \langle G'^{\sqcup\!\sqcup}(\mathbf{2}), \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \, dX_t \\
&= \int_0^S \langle \pi_{\leq N}(\exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1})) \sqcup\!\sqcup G'^{\sqcup\!\sqcup}(\mathbf{2}), \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \, dX_t \\
&= \int_0^S \langle \pi_{\leq N+\deg(G')}(\exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1}) \sqcup\!\sqcup G'^{\sqcup\!\sqcup}(\mathbf{2})), \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \, dX_t \\
&= \langle (\exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1}) \sqcup\!\sqcup G'^{\sqcup\!\sqcup}(\mathbf{2}))\mathbf{2}, \widehat{\mathbb{X}}_{0,S}^{\leq N+\deg(G')+1} \rangle.
\end{aligned}
$$

Similarly, since V is a polynomial,

$$
\int_0^S \langle \exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N} \rangle L(X_t) \, dt = \langle (\exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1}) \sqcup\!\sqcup L^{\sqcup\!\sqcup}(\mathbf{2}))\mathbf{1}, \widehat{\mathbb{X}}_{0,S}^{\leq N+\deg(V)+1} \rangle.
$$

Taking expectation, we obtain

$$
\begin{aligned}
\mathbb{E}\left[\int_0^S \langle \exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{\leq N} \rangle \, dY_t \right] &= \langle (\exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1}) \sqcup\!\sqcup G'^{\sqcup\!\sqcup}(\mathbf{2}))\mathbf{2}, \mathbb{E}[\widehat{\mathbb{X}}_{0,S}^{\leq N+\deg(G')+1}] \rangle \\
&\quad + \langle (\exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)\mathbf{1}) \sqcup\!\sqcup L^{\sqcup\!\sqcup}(\mathbf{2}))\mathbf{1}, \mathbb{E}[\widehat{\mathbb{X}}_{0,S}^{\leq N+\deg(V)+1}] \rangle.
\end{aligned}
$$

Using Theorem 6.5, we can deduce the result. $\square$

Remark 6.7. Once a Monte Carlo simulation of the expected signature of the rough path is generated, the optimization problems corresponding to the suprema in Corollary 6.6 are *purely deterministic*. Indeed, the objective function that needs to be maximized is a polynomial in the coefficients of the word l. In Section C.3 we discuss the implementation details for this step.

Remark 6.8. When the payoff process Y is a martingale then it follows from Doob's optional sampling theorem that optimal stopping value is zero. It is instructive to demonstrate how this fact can be observed in the form of the optional stopping problem in Corollary 6.6. When Y is a continuous martingale then it also admits a lift to a geometric rough path and therefore we may assume that $X = Y$. Now let w be an arbitrary word, then it holds

$$
\langle w\mathbf{12}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle = \int_0^t \left(\int_0^s \langle w, \widehat{\mathbb{X}}_{0,u}^{<\infty} \rangle du \right) \circ dX_s = \int_0^t \left(\int_0^s \langle w, \widehat{\mathbb{X}}_{0,u}^{<\infty} \rangle du \right) dX_s,
$$

where we have used that the Stratonovich-Itō correction is zero due to the finite variation of the integrand. By the martingale property of the integral in the above right hand side it follows that $\langle w\mathbf{12}, \mathbb{E}[\widehat{\mathbb{X}}_{0,S}^{\leq N}] \rangle = 0$ for all $N \geq 1$. Therefore, we see that all terms in the supremum in the right hand side of (6.6) vanish.

Remark 6.9. Note that similar formulas are available whenever Y is roughly given as a polynomial of the signature. We restrict ourselves to a representative class of examples below. We note here that payoffs of American options usually cannot be exactly represented in such a way. In particular, for the standard American (put) option, we have $Y_t = (K - X_t)_+$ for some $K > 0$, where X denotes the underlying asset price process. If we want to price American options using signature stopping policies, we have two possible remedies. We can approximate the payoff function by polynomials, which would allow us to directly apply Corollary 6.6. Alternatively, we can attach Y to the path

X, i.e., consider $\widetilde{X}_t \equiv (t, X_t, Y_t)$. Then, the corollary applies trivially, but at the price of increasing the dimension of the state space. The same strategy also works for more complicated functionals of the rough path $\widehat{\mathbb{X}}$. For instance, Y can be of the form $Y_t = g(t, \widetilde{Y}_t)$ where $\widetilde{Y}$ solves a rough differential equation

$$d\widetilde{Y}_t = b(\widetilde{Y}_t)\, dt + \sigma(\widetilde{Y}_t)\, d\mathbb{X}_t.$$

If g is sufficiently smooth, Y is *controlled by* $\mathbb{X}$ (cf. [FH14]) which guarantees that $\widetilde{X}_t \equiv (t, X_t, Y_t)$ can be lifted to a rough paths valued process.

7. Deep signature stopping policies

In the two preceding sections we focused on stopping policies θ that are linear functionals of the signature. Proposition 5.3 shows that this class of policies is sufficient to approximate a optimal stopping policy. Regarding numerical approximation and in order to obtain reasonable results it might however be necessary to step deeply into the signature, that is to approximate the signature upto a high truncation level. In this section we propose a different model for the stopping policy, in which the signature serves as an input feature to a deep neural network. Proposition 4.4 from above is the key link for defining a smooth loss-function for these models, which will be discussed in more detail in Section C.4. In this framework, the shuffle property of the signature is not directly needed and therefore it is reasonable to replace the signature by the log-signature, which allows for a low dimensional representation as explained below.

We will again assume throughout the section that $\mathbb{X} \in \Omega_T^p$ and that $Y = (Y_t)_{0 \le t \le T}$ is a real-valued continuous process adapted to the filtration generated by $\mathbb{X}$.

Definition 7.1 (Log-signature). We define the log-signature $\mathbb{L}^{<\infty} : [0, T] \to T((V))$ by

$$\mathbb{L}_t^{<\infty} = \log^{\otimes}(\mathbb{X}_t^{<\infty}), \quad 0 \le t \le T,$$

where $\log_{\otimes}$ is the tensor logarithm defined by

$$\log^{\otimes}(\mathbf{1} + \mathbf{a}) = \sum_{k=0}^{\infty} \frac{(-1)^{k+1}}{k} \mathbf{a}^{\otimes k} \in T((V)), \quad \mathbf{a} \in T((V)), \ \langle \varnothing, \mathbf{a} \rangle = 0.$$

We will also write $\widehat{\mathbb{L}}^{<\infty} = \log^{\otimes}(\widehat{\mathbb{X}}^{<\infty})$ for the log-signature of the time-augmented path $\widehat{\mathbb{X}}$. Recall that the signature is an element of the group $G(V)$. Therefore, the log-signature takes values in the set

$$\mathfrak{g}(V) := \log^{\otimes}(G(V)) = \left\{ \log^{\otimes}(\mathbf{g}) \mid \mathbf{g} \in G(V) \right\} \subset T((V)),$$

which forms a Lie-algebra with the bracket given by the commutator

$$[\mathbf{a}, \mathbf{b}] = \mathbf{a} \otimes \mathbf{b} - \mathbf{b} \otimes \mathbf{a}.$$

In fact, $\mathfrak{g}(V)$ is the free Lie-algebra over V (see e.g. [FV10, Section 7.3]) and therefore the log-signature can be expressed in terms of iterated Lie brackets of the basis elements $\{e_1, \cdots, e_d\}$, i.e. there exists explicit coefficients $(\lambda_{i_1,\ldots,i_n}(t))$ such that

$$\mathbb{L}_{0,t}^{<\infty} = \sum_{n=1}^{\infty} \sum_{i_1,\ldots,i_n=1}^{d} \lambda_{i_1,\ldots,i_n}(t) [e_{i_1}, [e_{i_2}, [\ldots, [e_{i_{n-1}}, e_{i_n}]\ldots]].$$

	$d = 1$		$d = 2$		$d = 3$		$d = 4$	
	$\sigma_{d,N}$	$\eta_{d,N}$	$\sigma_{d,N}$	$\eta_{d,N}$	$\sigma_{d,N}$	$\eta_{d,N}$	$\sigma_{d,N}$	$\eta_{d,N}$
$N = 1$	1	1	2	2	3	3	4	4
$N = 2$	2	1	6	3	12	6	20	10
$N = 3$	3	1	14	5	39	14	84	30
$N = 4$	4	1	30	8	120	32	340	90
$N = 5$	5	1	62	14	363	80	1364	294
$N = 6$	6	1	126	23	1092	196	5460	964

TABLE 1. Comparisons of the dimensionality of the representations of the truncated signature $\sigma_{d,N}$ and truncated log-signature $\eta_{d,N}$. (Values taken from [Rei17].)

This result already goes back to [Che57]. However, the above representation is not very efficient, due to linear dependencies (e.g. $[e_1, e_2] = -[e_2, e_1]$). An efficient basis is given by the *Lyndon basis* or more generally by a *Hall basis*. We refer to [Reu93] for a construction of such a basis. Define

$$\mathfrak{g}^N(V) = \pi_{\leq N}(\mathfrak{g}(V)) = \log^{\otimes}(G^N(V)),$$

the free step-N nilpotent Lie-algebra, then we have the following

Proposition 7.2. *For $N \in \{1, 2, \dots\}$ the dimension of $\mathfrak{g}^N(\mathbb{R}^d)$ as a linear vector space is given by*

$$(7.1) \qquad\qquad \eta_{d,N} := \sum_{n=1}^{N} \frac{1}{n} \sum_{k \mid n} \mu(n/k) d^k,$$

where the inner sum is taken over all divisors k of n and μ is the Möbius function.

Proof. The statement directly follows from the existence of a Hall basis of the free Lie-algebra and its dimensionality, which can be found in [Reu93, Theorem 6]. $\qquad\square$

Note that dimension of the truncated tensor algebra $T^N(\mathbb{R}^d)$ is given by $(d^{N+1} - 1)/(d - 1)$. Since the truncated signature always starts with $\mathbf{1}$, its representation in the linear space is effectively of dimension $\sigma_{d,N} = (d^{N+1} - d)/(d - 1)$. In Table 1 we compare the values $\sigma_{d,N}$ and $\eta_{d,N}$ for different values of d and N. While both values have the same asymptotic growth when $N \to \infty$, the efficient representation of the log-signature allows for numerical tractability up to higher truncations levels N. Finally note that the Baker-Campbell-Hausdorff formula, and therefore the log-signature of a piecewise linear path, can be directly calculated in a Hall basis. For an algorithm see for instance [CM08] and the summarizing note [Rei17].

We are now ready to present our class of deep signature stopping policies.

Definition 7.3. Define $\mathcal{T}_{\log} \subset \mathcal{T}$ to be the set of continuous stopping policies $\theta : \Lambda_T \to \mathbb{R}$ of the form

$$\theta(\widehat{\mathbb{X}}|_{[0,t]}) = \left(\theta_{\log} \circ \log^{\otimes}\right)(\widehat{\mathbb{X}}_{0,t}^{\leq N}),$$

for some $N \in \{1, 2, \dots\}$ and where $\theta_{\log} : \mathfrak{g}^N(\mathbb{R}^{d+1}) \to \mathbb{R}$ is a deep neural network of the form

$$(7.2) \qquad\qquad \theta_{\log} = A_0 \circ \varphi \circ A_1 \circ \varphi \circ \cdots \circ A_I,$$

where $A_I : \mathfrak{g}^N(\mathbb{R}^{d+1}) \to \mathbb{R}^q$, $A_i : \mathbb{R}^q \to \mathbb{R}^q$, $1 \leq i < I$, $A_0 : \mathbb{R}^q \to \mathbb{R}$, are affine maps for some $q, I \in \{1, 2, \dots\}$ and φ is an activation function, i.e. continuous and not a polynomial (e.g. $\varphi(x)_i = \max\{x_i, 0\}$ for $1 \leq i \leq q$ and $x \in \mathbb{R}^q$).

The following proposition, which is a consequence of Proposition 4.2, Proposition 5.3 and the universal approximation theorem for neural networks, states that the optimization over stopping policies in the class $\mathcal{T}_{\log}$ is sufficient.

Proposition 7.4. *Given* $\mathbb{E}[\|Y\|_\infty] < \infty$, *we have*

$$\sup_{\theta \in \mathcal{T}_{\log}} \mathbb{E}[Y_{\tau_\theta^r \wedge T}] = \sup_{\tau \in \mathcal{S}} \mathbb{E}[Y_{\tau \wedge T}].$$

Proof. Since $\mathcal{T}_{\log} \subset \mathcal{T}$ it follows from Proposition 4.2 and 5.3 that it suffices to show that

$$\sup_{\theta \in \mathcal{T}_{\text{sig}}} \mathbb{E}[Y_{\tau_\theta^r \wedge T}] \leq \sup_{\theta \in \mathcal{T}_{\log}} \mathbb{E}[Y_{\tau_\theta^r \wedge T}].$$

Therefore, let $\theta \in \mathcal{T}_{\text{sig}}$, i.e. there exists some $l \in T((\mathbb{R}^{d+1})^*)$ such that

$$\theta(\mathbb{X}|_{[0,t]}) = \langle l, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle = \langle l, \widehat{\mathbb{X}}_{0,t}^{\leq N} \rangle,$$

where $N := \max\{\deg(l), [p]\}$. We equip $G^N(\mathbb{R}^{d+1}) \subset T^N(\mathbb{R}^{d+1})$ with the subspace topology. Note that the inverse of the tensor logarithm $\log^\otimes$ is the tensor exponential map $\exp^\otimes$, which is defined analogously by its power series expansion (see also Example 2.3). Now define the continuous map

$$\psi : \mathfrak{g}^N(\mathbb{R}^{d+1}) \to \mathbb{R}, \quad \mathbf{x} \mapsto \langle l, \exp^\otimes(\mathbf{x}) \rangle$$

Let $K \subset G^N(\mathbb{R}^{d+1})$ be an arbitrary compact set and define $K' := \log^\otimes(K) \subset \mathfrak{g}^N(\mathbb{R}^{d+1})$, which is also compact. From [LLPS93, Theorem 1] it follows that there exists as sequence of functions $(\psi_n)_{n \geq 1}$ of the form (7.2) that converge uniformly on K' towards ψ. Hence, it also follows that the sequence of maps $(\psi_n \circ \log^\otimes)_{n \geq 1}$ converge uniformly on K towards the map $\langle l, \cdot \rangle : G^N(\mathbb{R}^{d+1}) \to \mathbb{R}$.

Next, note that the Lyons-lift map

$$\Lambda_T \to G^N(\mathbb{R}^{d+1}), \quad \widehat{\mathbb{X}}|_{[0,t]} \mapsto \widehat{\mathbb{X}}_{0,t}^{\leq N}$$

is continuous (see [FV10, Section 9.1]) and therefore it follows from above that for any compact set $\mathcal{K} \subset \widehat{\Omega}_T^p$, there exists a sequence $(\theta_n)_{n=1}^\infty \subset \mathcal{T}_{\log}$ such that

$$\lim_{n \to \infty} \sup_{\mathbb{X} \in \mathcal{K}, t \in [0,t]} |\theta_n(\mathbb{X}|_{[0,t]}) - \theta(\mathbb{X}|_{[0,t]})| = 0.$$

The rest of the proof is now completely analogous to proof of Proposition 5.3. $\square$

8. Numerical Examples

In this section we are going to evaluate the performance of the signature stopping methodology for the numerical approximation of the optimal stopping value. We have particularly chosen examples in which the underlying rough path is non-Markov. The technical steps that are necessary to implement a Monte-Carlo simulation are discussed in detail in Appendix C.

8.1. **Optimal stopping of fractional Brownian motion.** Let X^H be a fractional Brownian motion with Hurst parameter $H \in (0,1]$ and recall from Example 2.4 that the time augmented process $\widehat{X}^H$ has a lift to a geometric rough path. Also recall that in the case $H = 0.5$ the process X^H is a standard Brownian motion and otherwise it is not a semimartingale and not a Markov process. We are considering the optimal stopping problem with the payoff given by the underlying process itself, i.e. $Y \equiv X^H$. The payoff clearly satisfies the condition $\mathbb{E}(\|Y\|_{\infty;[0,T]}) < \infty$ for any $T < \infty$ and therefore our approximation results apply. We set $T = 1$ and use a time discretization with $J = 100$ steps for Monte-Carlo simulation of the payoff and the signature (see Section C.1 for more detail).

This example was recently studied in [BCJ19] and we will use the there presented values as benchmarks. The methodology in [BCJ19] was developed for Markov processes and is based on a parametrization of the optimal stopping policy by deep neural networks, one for each time in the discretization grid. The networks are then trained in a backwards induction using the dynamic programming principle. The methodology was applied to the fractional Brownian motion example by lifting the time-discretized process to a 100-dimensional Markov process including the entire past of the process.

Theoretical observations. Before evaluating the numerical results we present a few theoretical observations about the value of the optimal optimal stopping problem of the fractional Brownian motion.

In the case $H = 0.5$, the payoff process Y is a standard Brownian motion and hence a martingale. Therefore it follows from Doob's optional sampling theorem that the true value of the optimal stopping problem is zero.

The inclusion of the case $H = 1$ in the definition of the fractional Brownian motion is a matter of convention, in any case, the right-hand side of (2.4) defines the positive semidefinite covariance kernel $\rho_1(s,t) = s \cdot t$ for $s,t \geq 0$. The process $X := (t \cdot \xi)_{t \geq 0}$ with a standard normal distributed random variable ξ admits this covariance structure. Since the sample paths of this process are given by straight lines, one can easily verify that an optimizing sequence of stopping times $(\tau_n)_{n \geq 1}$ is given by

$$\tau_n = \begin{cases} T, & X_{1/n} > 0 \\ 1/n, & X_{1/n} \leq 0 \end{cases}, \quad n = 0, 1, 2, \dots$$

The optimal stopping value for $H = 1$ is therefore given by $\lim_{n \to \infty} \mathbb{E}[Y_{\tau_n}] = \mathbb{E}[T \cdot (\xi)_+] = T/\sqrt{2\pi}$. Furthermore, we also see that the value of the time discretized problem is given by $\mathbb{E}[Y_{\tau_J}] = 0.99/\sqrt{2\pi} \approx 0.395$ for $J = 100$ and $T = 1$. We point out that this discussion of the case $H = 1$ was already presented in [BCJ19].

On the other side of the spectrum, that is when H approaches 0, the covariance kernel of the fractional Brownian motion (2.4) converges pointwise to the following kernel

$$(8.1) \qquad \rho_0(s,t) = \begin{cases} 0, & t = 0 \vee s = 0. \\ 1, & t = s \wedge s > 0, \\ \frac{1}{2}, & \text{else.} \end{cases} \quad , \quad t, s \geq 0.$$

The Kolmogorov extension theorem implies that a Gaussian process with this covariance structure exist, however such a process is not measurable in the usual sense of stochastic processes (see [Sto14, Section 19.5]) and therefore the associated optimal stopping problem is not well defined. See also

[FKS16], [NR18], [HN20] and [BHP20] for regularizations and normalizations of the fractional Brownian motion and the convergence as H approaches 0. For the comparison with numerical approximations it is nevertheless interesting to study the discrete time optimal stopping problem of a process with the covariance structure ρ_0. Let $(\xi_0, \xi_1, \ldots, \xi_J)$ be a finite sequence of independent standard normal distributed random variables and note that the process

$$\widetilde{Y}_j := (\xi_j - \xi_0)/\sqrt{2}, \quad j = 0, 1, \ldots, J,$$

has the covariance structure

$$\mathbb{E}\big[\widetilde{Y}_i \widetilde{Y}_j\big] = \rho_0(i,j) = \rho_0(t_i, t_j), \quad i, j \in \{0, 1, \ldots, J\}, \ \ t_j = Tj/J.$$

In Appendix B we present an explicit solution to the discrete time optimal stopping problem associated to the process $\widetilde{Y}$. For $J = 100$ the optimal stopping value is approximately given by 1.5830 and for $J \to \infty$ the value converges to infinity (see Proposition B.3 and Remark B.4). At this point, let us also mention that the value of the discrete time stopping problem depends continuously on the Hurst parameter $H \in [0, 1]$. Indeed, one has to verify that the distribution of the process $(X_{t_j}^H)_{j=0,1,\ldots,J}$ depends continuously on H in the *weak adapted topology* (see [BBBE20]). However, this is a consequence of the fact that the process is Gaussian and the densities of the law (receptively conditional laws) are everywhere continuous in H. Since the value of the discrete time optimal stopping problem is a lower-bound to the value of the continuous time problem we have the following

Proposition 8.1. *The value of the optimal stopping problem of the fractional Brownian motion* $(X_t^H)_{0 \leq t \leq T}$ *diverges as $H \to 0$, more precisely we have*

$$\lim_{H \to 0} \sup_{\tau \in \mathcal{S}} \mathbb{E}\left[X_{\tau \wedge T}^H\right] = \infty.$$

Evaluation of the numerical results. We have approximated the value of the optimal stopping problem using both linear signature stopping policies ($\mathcal{T}_{\mathrm{sig}}$) and deep signature stopping policies ($\mathcal{T}_{\log}$), for Hurst parameters $H \in \{0.1, 0.2, \ldots, 1.0\}$ and $H = 0$ (in the sense of the theoretical observations above). The resulting low-biased estimates are presented in Table 2 (see Section C.5 for a description of the calculation of low-biased estimates). The values associated to the linear stopping policy presented in Table 2 are obtained using a stochastic gradient decent method based one the form of the expected payoff given in Proposition 4.4. The technical details for this procedure can be found in Section C.4. We have used $2^{19} = 524,288$ samples for the optimization and $2^{23} = 8,388,608$ samples for the resimulation. The values associated to the deep stopping policy were obtained using a network architecture consisting of $I = 2$ hidden layers, $q = \eta_{2,N} + 30$ neurons for each layer, depending on the truncation level N of the log-signature, and relu activation function. The networks were trained with $2^{21} = 2,097,152$ samples and the loss-function presented in Section C.4. For the resimulation we used again 2^{23} samples.

Comparing the values in Table 2, it is not very surprising that we observe a worse performance of the linear method compared with the non-linear method using the same truncation level. The values obtained with the deep signature policies achieve the exact values for $H = 1$ and $H = 0$ (see the theoretical observations above) and come reasonably close to he low-biased values presented [BCJ19], already for a truncation level $N = 2$. This observation becomes more impressing when we recall from Tabel 1 that the dimension of the log-signature for $d = 2$ and $N = 2, 3, 4, 5$ is given by $\eta_{2,N} = 3, 5, 8, 14$. These dimensions must be seen in contrast to the 100-dimensional input features

H	lin. signature policies			deep signature policies					BCJ	
	$N=3$	$N=4$	$N=5$	$N=1$	$N=2$	$N=3$	$N=4$	$N=5$	low	up
0.0	1.161	1.331	1.430	1.384	1.582	1.583	1.583	1.583	-	-
0.1	0.760	0.858	0.916	0.929	1.030	1.038	1.042	1.043	1.048	1.049
0.2	0.468	0.523	0.553	0.584	0.637	0.646	0.649	0.651	0.658	0.659
0.3	0.257	0.285	0.293	0.329	0.355	0.361	0.363	0.364	0.369	0.380
0.4	0.107	0.116	0.118	0.140	0.149	0.152	0.152	0.153	0.155	0.158
0.5	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.005
0.6	0.081	0.084	0.085	0.106	0.112	0.113	0.114	0.114	0.115	0.118
0.7	0.146	0.152	0.154	0.191	0.198	0.201	0.201	0.202	0.206	0.207
0.8	0.199	0.207	0.208	0.263	0.269	0.272	0.273	0.273	0.276	0.278
0.9	0.247	0.256	0.258	0.328	0.332	0.334	0.334	0.334	0.336	0.339
1.0	0.296	0.307	0.309	0.395	0.395	0.395	0.395	0.395	0.395	0.395

TABLE 2. Low-biased estimates to the optimal stopping value of the fractional Brownian motion for different Hurst parameters H, obtained with linear ($\mathcal{T}_{\mathrm{sig}}$) and deep ($\mathcal{T}_{\log}$) signature stopping policies. N is the truncation level of the log-/signature. For the deep signature stopping policies we used neural networks with $I = 2$ hidden layers and $q = \eta_{2,N} + 30$ neurons. The values in the two rightmost columns are the low- and upper-biased estimates obtained in [BCJ19]. The overall Monte-Carlo error in the resimulation is below 0.0004. The definition of the discrete time optimal stopping problem corresponding to the case $H = 0.0$ was discussed in the theoretical observations above.

used in [BCJ19]. This suggests that also in the context of optimal stopping of non-Markov processes the log-signature can serve as an efficient compression of the path.

Figure 1 visualizes low-biased estimates to the value of optimal stopping problem of the fractional Brownian motion using time discretizations with $J = 100$ and $J = 1000$ steps. The values were obtained using a deep signature stopping policy with a signature truncation level $N = 3$ and the network configuration $I = 2$, $q = 35$. Note that the dimension of the (log-)signature obviously does not depend on the size of the time discretization grid. Therefore, the number of model parameters of the signature stopping policy does not change when increasing J. We see that the approximate values for $H \ll 0.5$ are significantly higher when using the finer discretization grid. An observation of this sort was already foreshadowed by the divergence of the optimal stopping value for $H \to 0$ from Proposition 8.1.

Linearization. Before closing this section, we are going to discuss the numerical application of the results from Section 6. Since the payoff coincides with the underlying path itself, Corollary 6.6 applies and after estimating the expected signature for some truncation level, this allows to approximate the value of the optimal stopping problem by a numerical approximation of the supremum in the right-hand side of (6.6). Section C.3 describes the technical details for this step. However, it turned out that the values obtain with this procedure are drastically worse then the values obtained with any of the other methods that we have studied. The difficulty lies in the tradeoff between the optimisation constrained $\{l \in T(V^*) \mid \deg(l) + |l| \leq K\}$ and the truncation level of the signature. More precisely, we understand from the estimate in Lemma 6.2 that choosing K large requires to also choose the truncation level N large in order to decrease the error that is introduced
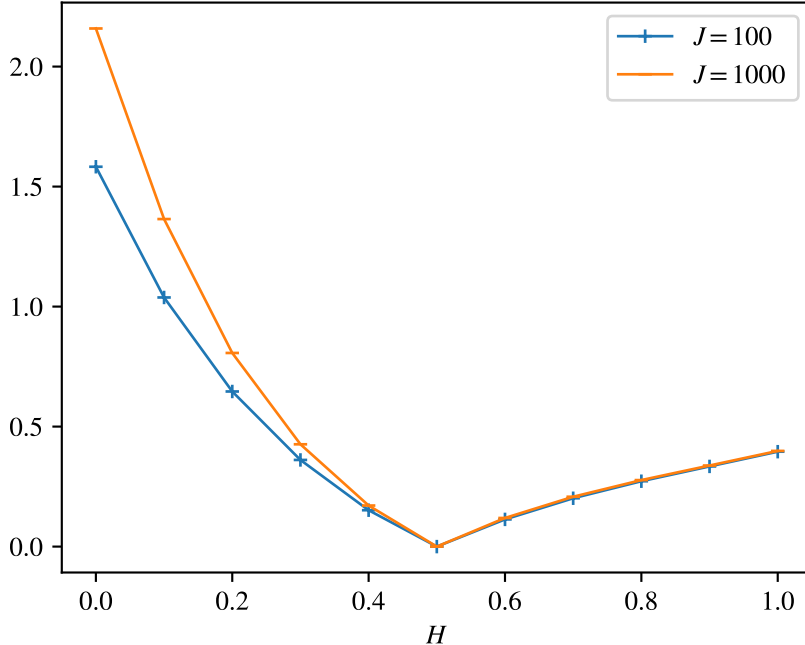
FIGURE 1. A visualization of lower biased estimates to the optimal stopping value of fractional Brownian motion for different Hurst parameters $H \in \{0.0, 0.1, \ldots, 1.0\}$, using time discretization grids with $J = 100$ and $J = 1000$ steps. The values were calculated with the deep signature stopping rule using a truncation level $N = 3$ for the log-signature and a neural network with $I = 2$ hidden layers and $q = 5 + 30$ nodes per layer.

from replacing the exponential with the shuffle exponential. However, increasing N also increases computational costs and, for this example, the limits of our computational resources were reached before obtaining reasonable results.

8.2. **An American put option in a rough electricity market.** We consider the following rough model for the spot electricity prices S proposed in [Ben17]

$$S_t = \exp(X_t) = \exp(\Lambda_t + Z_t), \quad 0 \le t \le T,$$

where lambda Λ is a seasonality component and Z is the rough base signal. Note that we have refrained from using a spike signal in the model for instructive reasons. Since we are interested in pricing an option on S and in this context prices are usually corrected for seasonality influences, we will assume without much loss of generality that

$$\Lambda \equiv x_0 \in \mathbb{R}_+.$$

Regarding the rough component, we follow the propositions in [Ben17] and finally choose the following model for the log-price

$$X_t = x_0 + \frac{\sigma}{c_\alpha} \int_{-\infty}^{t} (t - s)^\alpha e^{-\lambda(t-s)} dW_s, \quad 0 \le t \le T,$$

K	$N = 1$	$N = 2$	$N = 3$	$N = 4$	EUR
80	3.923	4.062	4.076	4.073	0.516
90	9.841	10.185	10.206	10.228	2.139
100	17.753	18.279	18.341	18.363	5.770
110	26.679	27.352	27.397	27.423	11.580
120	36.099	36.854	36.945	36.898	19.128

TABLE 3. Low-biased estimates of the value of the American put option with $T = 1.0$ and different strikes K on the energy spot price with model parameters $x_0 = 100$, $\alpha = -0.4$, $\lambda = 0.02$, $\sigma = 0.2$ and $r = 0.05$. A time discretization grid with $J = 100$ points was used. The values were obtained with deep signature stopping policies using $I = 2$ hidden layer and $q = \eta_N + 30$ neurons per layer. The rightmost column represents the prices of the European option. The overall Monte Carlo error in the resimulation is below 0.0005.

where $\alpha \in (-1/2, 1/2)$, $\lambda, \sigma > 0$, c_α is a normalizations constant such that $\mathbb{E}(X_t^2) = \sigma^2$ and W is a two sided standard Brownian motion. Note that X is a stationary Gaussian process and has the following auto-correlation function (see [BN12])

$$\mathrm{Corr}(X_t, X_{t+h}) = \frac{2^{-\alpha + \frac{1}{2}}}{\Gamma(\alpha + \frac{1}{2})} (\lambda h)^{\alpha + \frac{1}{2}} \widetilde{K}_{\alpha + \frac{1}{2}}(\lambda h)$$

where $\widetilde{K}_\nu(x) = e^{-x} K_\nu(x)$ and K_ν is the modified Bessel function of the second kind. One can easily verify that the paths of X are γ-Hölder continuous for any $\gamma < \alpha + 1/2$ (see e.g. [Ben20]). Hence, by the same arguments as in Example 2.4, it follows that X and the time augmented process $\widehat{X}$ have a lift to a geometric rough path.

We are interested in approximating the price of an American put option on S, which is given by

$$\sup_{\tau \in \mathcal{S}} \mathbb{E}\left[e^{-r\tau} (K - S_{\tau \wedge T})_+ \right],$$

where $K > 0$ is the strike and $r \geq 0$ is an interest rate. Hence, the payoff process of the corresponding optimal stopping problem is a continuous functional of underlying path given by $Y_t = e^{-rt}(K - \exp(X_t))_+$.

We set $T = 1$ and for Monte-Carlo simulation use a time discretization with $J = 100$ steps. Note that the sample trajectories of X can be obtained using the exact form of the above autocorrelation function. In order to obtain simulations with a constant initial value, we sampled from the conditional law $\mathbb{P}(\cdot | X_0 = x_0)$. For example, when generating the samples of X with the Cholesky decomposition of the covariance matrix, this is simply achieved by setting the first component of the standard normal input noise to zero.

In Table 3 we present low-biased estimates to the value of American put option obtained with deep signature stopping policies. As model parameters we have chosen $\alpha = -0.4$, $\lambda = 0.02$ (these values are based on statistical estimates given in [Ben17]), initial value $x_0 = 100$, volatility $\sigma = 0.2$ and interest rate $r = 0.05$. The network architecture for the deep stopping policy consisted of $I = 2$ hidden layers, $q = \eta_{2,N} + 30$ neurons per layer and relu activation functions. The networks were trained using 2^{20} sample paths and the low-biased estimates were calculated with 2^{23} sample paths. As a reference value we have also presented an estimate of the price of the European put option,

which is given by $\mathbb{E}[e^{-rT}(K - S_T)_+]$. We observe that the values significantly improve when moving from truncation level $N = 1$ to $N = 2$. For higher levels this improvement quickly saturates, which suggests that the values are approaching the true value of the option.

References

[AGG+18] Imanol Perez Arribas, Guy M Goodwin, John R Geddes, Terry Lyons, and Kate EA Saunders. A signature-based machine learning model for distinguishing bipolar disorder and borderline personality disorder. *Translational psychiatry*, 8(1):1–7, 2018.

[BBBE20] Julio Backhoff-Veraguas, Daniel Bartl, Mathias Beiglböck, and Manu Eder. All adapted topologies are equal. *Probability Theory and Related Fields*, 178(3):1125–1172, 12 2020.

[BCJ19] Sebastian Becker, Patrick Cheridito, and Arnulf Jentzen. Deep optimal stopping. *Journal of Machine Learning Research*, 20(74):1–25, 2019.

[Ben17] Mikkel Bennedsen. A rough multi-factor model of electricity spot prices. *Energy Economics*, 63:301–313, March 2017.

[Ben20] Mikkel Bennedsen. Semiparametric estimation and inference on the fractal index of gaussian and conditionally gaussian time series data. *Econometric Reviews*, 39(9):875–903, 2020.

[BHP20] Christian Bayer, Fabian Harang, and Paolo Pigato. Log-Modulated Rough Stochastic Volatility Models. 2020.

[BN12] Ole E. Barndorff-Nielsen. Notes on the gamma kernel. Technical Report 03, T.N. Thiele Centre, Department of Mathematics, Aarhus University, 2012.

[BTW20] Christian Bayer, Raúl Tempone, and Sören Wolfers. Pricing american options by exercise rate optimization. *Quantitative Finance*, pages 1–12, 2020.

[CF10] Rama Cont and David Fournie. A functional extension of the ito formula. *Comptes Rendus Mathematique*, 348(1-2):57–61, 2010.

[Che57] Kuo-Tsai Chen. Integration of paths, geometric invariants and a generalized baker- hausdorff formula. *Annals of Mathematics*, 65(1):163–178, 1957.

[CL16] Ilya Chevyrev and Terry Lyons. Characteristic functions of measures on geometric rough paths. *The Annals of Probability*, 44(6):4049–4082, 2016.

[CM08] Fernando Casas and Ander Murua. An efficient algorithm for computing the baker-campbell-hausdorff series and some of its applications. *Journal of Mathematical Physics*, 50, 11 2008.

[Dup19] Bruno Dupire. Functional Itô calculus. *Quant. Finance*, 19(5):721–729, 2019.

[Faw02] Thomas Fawcett. *Problems in stochastic analysis : connections between rough paths and non-commutative harmonic analysis.* PhD thesis, University of Oxford, 2002.

[FH14] Peter K. Friz and Martin Hairer. *A Course on Rough Paths with an introduction to regularity structures*, volume XIV of *Universitext*. Springer, Berlin, 2014.

[FKS16] Y. V. Fyodorov, B. A. Khoruzhenko, and N. J. Simm. Fractional Brownian motion with Hurst index $H = 0$ and the Gaussian Unitary Ensemble. *The Annals of Probability*, 44(4):2980 – 3031, 2016.

[FV10] Peter K. Friz and Nicolas B. Victoir. *Multidimensional stochastic processes as rough paths*, volume 120 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2010. Theory and applications.

[HL10] Ben Hambly and Terry Lyons. Uniqueness for the signature of a path of bounded variation and the reduced path group. *Annals of Mathematics*, pages 109–167, 2010.

[HN20] Paul Hager and Eyal Neuman. The Multiplicative Chaos of $H = 0$ Fractional Brownian Fields. *arXiv e-prints*, page arXiv:2008.01385, August 2020.

[KLPA20] Jasdeep Kalsi, Terry Lyons, and Imanol Perez Arribas. Optimal execution with rough path signatures. *SIAM J. Financial Math.*, 11(2):470–493, 2020.

[LCL07] Terry J. Lyons, Michael Caruana, and Thierry Lévy. *Differential equations driven by rough paths*, volume 1908 of *Lecture Notes in Mathematics*. Springer, Berlin, 2007. Lectures from the 34th Summer School on Probability Theory held in Saint-Flour, July 6–24, 2004, With an introduction concerning the Summer School by Jean Picard.

[LLPS93] Moshe Leshno, Vladimir Ya. Lin, Allan Pinkus, and Shimon Schocken. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural Networks*, 6(6):861–867, 1993.

[NR18]   Eyal Neuman and Mathieu Rosenbaum. Fractional Brownian motion with zero Hurst parameter: a rough volatility viewpoint. *Electronic Communications in Probability*, 23(none):1 – 12, 2018.

[Rei17]   Jeremy Reizenstein. Calculation of iterated-integral signatures and log signatures, 2017.

[Reu93]   Christophe Reutenauer. *Free Lie algebras*, volume 7 of *London Mathematical Society Monographs. New Series*. The Clarendon Press Oxford University Press, New York, 1993. Oxford Science Publications.

[RG20]   Jeremy F Reizenstein and Benjamin Graham. Algorithm 1004: The iisignature library: Efficient calculation of iterated-integral signatures and log signatures. *ACM Transactions on Mathematical Software (TOMS)*, 46(1):1–21, 2020.

[Shi08]   A. N. Shiryaev. *Optimal stopping rules*, volume 8 of *Stochastic Modelling and Applied Probability*. Springer-Verlag, Berlin, 2008. Translated from the 1976 Russian second edition by A. B. Aries, Reprint of the 1978 translation.

[Sto14]   J.M. Stoyanov. *Counterexamples in Probability: Third Edition*. Dover Books on Mathematics. Dover Publications, 2014.

[Wiś94]   Andrzej Wiśniewski. The structure of measurable mappings on metric spaces. *Proc. Amer. Math. Soc.*, 122(1):147–150, 1994.

Appendix A. Technical aspects of stopped rough paths

Recall from Definition 3.1 the space stopped rough paths Λ_T and its metric. The following lemma characterizes the topology on Λ_T.

Lemma A.1. *The topology on Λ_T coincides with the final topology induced by the map $\varphi \colon [0,T] \times \widehat{\Omega}_T^p \to \Lambda_T$, $\varphi(t, \widehat{\mathbb{X}}) = \widehat{\mathbb{X}}|_{[0,t]}$. Moreover, Λ_T is Polish.*

Proof. A set $U \subset \Lambda_T$ is open with respect to the final topology if and only if $\varphi^{-1}(U)$ is open in $[0,T] \times \widehat{\Omega}_T^p$. One can easily check that φ is continuous for the topology induced by d, therefore $\varphi^{-1}(U)$ is open for every open set $U \subset \Lambda_T$. Now assume that $\varphi^{-1}(U)$ is open for a set $U \subset \Lambda_T$. Let $\mathbb{X}|_{[0,t]} \in U$ and choose $\mathbb{Y}|_{[0,s]} \in \Lambda_T$ with $d(\mathbb{X}|_{[0,t]}, \mathbb{Y}|_{[0,s]}) < \varepsilon$. Our goal is to prove that $\mathbb{Y}|_{[0,s]} \in U$ for ε chosen sufficiently small. Note that

$$\varphi^{-1}(\mathbb{X}|_{[0,t]}) = \{(t, \widetilde{\mathbb{X}}) \, : \, \widetilde{\mathbb{X}}|_{[0,t]} = \mathbb{X}|_{[0,t]}\}.$$

Assume $s \geq t$ first. Then $d(\mathbb{X}|_{[0,t]}, \mathbb{Y}|_{[0,s]}) < \varepsilon$ implies that

$$|t - s| < \varepsilon \quad \text{and} \quad d_{p-\mathrm{var};[0,s]}(\widetilde{\mathbb{X}}|_{[0,s]}, \mathbb{Y}|_{[0,s]}) < \varepsilon$$

where $\widetilde{\mathbb{X}}|_{[0,s]}$ is the stopped path defined on $[0,s]$ as explained in Definition 3.1. Let $\widetilde{\mathbb{X}} = \widetilde{\mathbb{X}}|_{[0,T]} \in \widehat{\Omega}_T^p$ be the stopped path defined on the whole time interval $[0,T]$. Since $(t, \widetilde{\mathbb{X}}) \in \varphi^{-1}(\mathbb{X}|_{[0,t]}) \subset \varphi^{-1}(U)$ and $\varphi^{-1}(U)$ is open, there is a $\delta > 0$ such that whenever $u \in (t - \delta, t + \delta)$ and $d_{p-\mathrm{var};[0,T]}(\widetilde{\mathbb{X}}, \widetilde{\mathbb{Y}}) < \delta$, we have $(u, \widetilde{\mathbb{Y}}) \in \varphi^{-1}(U)$. Choosing ε sufficiently small, we can assume that $s \in (t - \delta, t + \delta)$. Define $\widetilde{\mathbb{Y}} = \widetilde{\mathbb{Y}}|_{[0,T]} \in \widehat{\Omega}_T^p$ as in Definition 3.1. Then $(s, \widetilde{\mathbb{Y}}) \in \varphi^{-1}(\mathbb{Y}|_{[0,s]})$ and

$$d_{p-\mathrm{var};[0,T]}(\widetilde{\mathbb{X}}, \widetilde{\mathbb{Y}}) \leq C_p(d_{p-\mathrm{var};[0,s]}(\widetilde{\mathbb{X}}, \widetilde{\mathbb{Y}}) + d_{p-\mathrm{var};[s,T]}(\widetilde{\mathbb{X}}, \widetilde{\mathbb{Y}}))$$
$$= C_p d_{p-\mathrm{var};[0,s]}(\widetilde{\mathbb{X}}|_{[0,s]}, \mathbb{Y}|_{[0,s]}) \leq C_p \varepsilon.$$

Choosing ε small, we conclude $(s, \widetilde{\mathbb{Y}}) \in \varphi^{-1}(U)$ and thus $\mathbb{Y}|_{[0,s]} \in U$. For $s \leq t$, we can argue similarly which proves that both topologies indeed coincide. Concerning the second statement, separability follows from the separability of $[0,T] \times \widehat{\Omega}_T^p$ and the fact that φ is a continuous surjection.

To prove that Λ_T is complete with respect to the metric d is straightforward and follows from the fact that $[0, T]$ and $\widehat{\Omega}_T^p$ are complete. $\qquad\square$

Corollary A.2. *Let Z be any topological space. A map $g\colon \Lambda_T \to Z$ is continuous if and only if the map $[0, T] \times \widehat{\Omega}_T^p \ni (t, \widehat{\mathbb{X}}) \mapsto g(\widehat{\mathbb{X}}|_{[0,t]}) \in Z$ is continuous.*

Proof. Follows from the universal property of the final topology. $\qquad\square$

Appendix B. Solution of an optimal stopping problem corresponding to a discrete time fractional Brownian motion with zero Hurst parameter

Recall from Remark 8.1 that, as H tends to zero, the covariance kernel of a fractional Brownian motion converges pointwise to the kernel ρ_0 given in (8.1). In this section we are going to derive an explicit solution to the optimal stopping problem associated to a discrete time Gaussian process with the covariance kernel ρ_0. Therefore let $(\xi_0, \xi_1, \ldots, \xi_J)$ be a finite sequence of independent standard normal distributed random variables defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Define the process

$$Y = (Y_j)_{j=0,1,\ldots,J}, \quad Y_j = \xi_j - \xi_0,$$

and note that its covariance structure is given by $\mathbb{E}[Y_i Y_j] = 2 \cdot \rho_0(i, j)$. For ease of notation we refrain from scaling the process by $1/\sqrt{2}$. Further note that the filtration that the process Y generates is given by

$$\mathcal{F}_0 = \{\emptyset, \Omega\}, \quad \mathcal{F}_j = \sigma(\xi_1 - \xi_0, \ldots, \xi_j - \xi_0).$$

We are interested in finding the solution to the discrete time optimal stopping problem associated to the payoff Y. Therefore denote by $\mathcal{S}_J$ the set of discrete stopping times with respect to the filtration $(\mathcal{F}_j)_{j=0,1,\ldots,J}$ and define the value process

$$V_j = \sup_{\sigma \in \mathcal{S}_J, \sigma \geq j} \mathbb{E}\left[Y_\sigma \Big| \mathcal{F}_j\right], \quad j = 0, 1, \ldots, J.$$

Then we have $V_J = Y_J$ and by the dynamic programming principle it follows

(B.1) $$V_j = \max\left\{Y_j, \ \mathbb{E}[V_{j+1}|\mathcal{F}_j]\right\}, \quad j = 0, 1, \ldots, J - 1.$$

We are now going to proceed by recursively calculating the conditional expectation in the dynamic programming principle. Therefore define

$$\overline{\xi_{1:j}} := \frac{1}{j} \sum_{k=1}^{j} \xi_k$$

and note that we have the following

Lemma B.1. *Let $j \in \{1, \ldots, J\}$, then the conditional expectation of ξ_0 given $\mathcal{F}_j$ is given by*

$$\mathbb{E}[\xi_0|\mathcal{F}_j] = \frac{j}{j+1}(\xi_0 - \overline{\xi_{1:j}}).$$

Proof. First note that the random variable in the above right-hand side is mean-zero Gaussian and is $\mathcal{F}_j$-measurable which can easily be seen from the identity

$$\frac{j}{j+1}(\xi_0 - \overline{\xi_{1:j}}) = \frac{1}{j+1}\sum_{i=1}^{j} Y_i.$$

On the other hand, we have for all $A \in \mathcal{F}_j$

$$\mathbb{E}\left[\left(\xi_0 - \frac{j}{j+1}(\xi_0 - \overline{\xi_{1:j}})\right)1_A\right] = \mathbb{E}\left[\frac{1}{j+1}(\xi_0 + \xi_1 + \cdots + \xi_j)1_A\right] = 0$$

since the random variable $\xi_0 + \xi_1 + \cdots + \xi_j$ is independent of $\mathcal{F}_j$. Indeed, every $Y_i = \xi_i - \xi_0$ for $i \le j$ is uncorrelated with $\xi_0 + \xi_1 + \cdots + \xi_j$, hence due to Gaussianity independent. □

Remark B.2. We can observe that the process $(Y_j, \xi_0 - \overline{\xi_{1:j}})_{j=0,1,\ldots,J}$ is a two-dimensional Markov process adapted to the filtration $(\mathcal{F}_j)_{j=0,1,\ldots,J}$. Indeed, recall that $Y_j = \xi_j - \xi_0$ and note that ξ_j is independent of $\mathcal{F}_i$ for all $i < j$. Further, we see that the conditional law of ξ_0 given $\mathcal{F}_i$ is Gaussian. From the above lemma it follows that the conditional mean is given by $\xi_0 - \overline{\xi_{1:i}}$ and following a simple Gaussian calculation we see that the conditional variance is deterministic (depending only on i and j).

The conditional expectation of the value function V_J with respect to $\mathcal{F}_{J-1}$ is then given by

$$\mathbb{E}[V_J|\mathcal{F}_{J-1}] = \mathbb{E}[-\xi_0|\mathcal{F}_{J-1}] = \frac{J-1}{J}(\overline{\xi_{1:J-1}} - \xi_0).$$

Hence, plugging into the Bellman equation (B.1) we obtain for the value process at $J-1$

$$V_{J-1} = \max\left\{\xi_{J-1} - \xi_0, \frac{J-1}{J}(\overline{\xi_{1:J-1}} - \xi_0)\right\}.$$

Using the dynamic programming principle we are now going to show the following

Proposition B.3. *For all $j \in \{1,\ldots,J-1\}$ it holds*

$$V_j = \max\left\{\xi_j - \xi_0, \mu_{j+1} + \frac{j}{j+1}(\overline{\xi_{1:j}} - \xi_0)\right\},$$

and $V_0 = \mu_1$, where

$$\mu_J = 0, \quad \mu_j = \sqrt{\frac{j}{j+1}}\gamma\left(\mu_{j+1}\sqrt{\frac{j+1}{j}}\right), \quad j = J-1, J-2, \ldots, 1$$

where $\gamma(x) = \phi(x) + x\Phi(x)$ and Φ and ϕ are the cdf and pdf of the standard normal distribution.

Remark B.4. Note that for any $j \in \{1,\ldots,J-1\}$ it holds $\mu_j \ge \gamma(2 \cdot \mu_{j+1})/2$. Hence it follows that $\mu_1 \ge f^{\circ(J-1)}(0)$ with $f(x) := \gamma(2x)/2$. Since f is continuous and $f(x) > x$ for all $x \in [0, \infty)$ it follows that $f^{\circ J}(0) \to \infty$ for $J \to \infty$. This confirms what one intuitively expects, that the value of the above optimal stopping problem converges to infinity as J tends to infinity.

Proof. The claim was already shown for $j = J - 1$. We are going to use a backwards induction to proof the claim for $0 \leq j < J - 1$. In order to explicitly calculate the conditional expectation of V_{j+1} given $\mathcal{F}_j$, we are going the derive a orthogonal decomposition of each of the random variables in the above maximum with respect to $\mathcal{F}_j$. For the first random variable we have

$$\xi_{j+1} - \xi_0 = \left(\xi_{j+1} - \frac{1}{j+1}\xi_0 - \frac{j}{j+1}\overline{\xi_{1:j}} \right) + \frac{j}{j+1}(\overline{\xi_{1:j}} - \xi_0)$$

and for the second we have

$$\begin{aligned}
\frac{j+1}{j+2}(\overline{\xi_{1:j+1}} - \xi_0) &= \frac{1}{j+2}\xi_{j+1} + \frac{j}{j+2}\overline{\xi_{1:j}} - \frac{j+1}{j+2}\xi_0 \\
&= \frac{1}{j+2}\xi_{j+1} + \frac{j}{j+2}(\overline{\xi_{1:j}} - \xi_0) - \frac{1}{j+2}\xi_0 \\
&= \frac{1}{j+2}\xi_{j+1} + \frac{j}{j+2}(\overline{\xi_{1:j}} - \xi_0) + \frac{1}{j+2}\frac{j}{j+1}(\overline{\xi_{1:j}} - \xi_0) \\
&\quad - \frac{1}{j+2}\left(\frac{1}{j+1}\xi_0 + \frac{j}{j+1}\overline{\xi_{1:j}} \right) \\
&= \frac{1}{j+2}\left(\xi_{j+1} - \frac{1}{j+1}\xi_0 - \frac{j}{j+1}\overline{\xi_{1:j}} \right) + \frac{j}{j+1}(\overline{\xi_{1:j}} - \xi_0).
\end{aligned}$$

For simplifying notation in what follows, we define

$$\eta_j := \sqrt{\frac{j+1}{j+2}}\left(\xi_{j+1} - \frac{1}{j+1}\xi_0 - \frac{j}{j+1}\overline{\xi_{1:j}} \right)$$

which is a standard Gaussian random variable orthogonal to $\mathcal{F}_j$, indeed note that

$$\mathbb{V}\left(\xi_{j+1} - \frac{1}{j+1}\xi_0 - \frac{j}{j+1}\overline{\xi_{1:j}} \right) = 1 + \frac{1}{(j+1)^2} + \frac{j}{(j+1)^2} = \frac{j+2}{j+1}.$$

Using this orthogonal decomposition and the induction claim we then have

$$\mathbb{E}[V_{j+1}|\mathcal{F}_j] = \sqrt{\frac{j+2}{j+1}}\mathbb{E}\left[\max\left\{ \eta_j, \ \mu_{j+2}\sqrt{\frac{j+1}{j+2}} + \frac{1}{j+2}\eta_j \right\} \right] + \frac{j}{j+1}(\overline{\xi_{1:j}} - \xi_0).$$

Resolving the above expectation is now simple one-variate Gaussian calculation. Therefore let $0 \leq \alpha < 1$, $\beta \in \mathbb{R}$, then note that

$$\begin{aligned}
\mathbb{E}[\max\{\eta_j, \beta + \alpha\eta_j\}] &= \int_{\beta/(1-\alpha)}^{\infty} x\phi(x)dx + \int_{-\infty}^{\beta/(1-\alpha)} (\beta + \alpha x)\phi(x)dx \\
&= \phi(\beta/(1-\alpha)) + \beta\Phi(\beta/(1-\alpha)) - \alpha\phi(\beta/(1-\alpha)) \\
&= (1-\alpha)\phi(\beta/(1-\alpha)) + \beta\Phi(\beta/(1-\alpha)) \\
&= (1-\alpha)\gamma(\beta/(1-\alpha)).
\end{aligned}$$

Therefore we finally have

$$\mathbb{E}[V_{j+1}|\mathcal{F}_j] = \sqrt{\frac{j+2}{j+1}}\frac{j+1}{j+2}\gamma\left(\mu_{j+2}\sqrt{\frac{j+1}{j+2}}\frac{j+2}{j+1} \right) + \frac{j}{j+1}(\overline{\xi_{1:j}} - \xi_0)$$

and the rest of the claim follows from the dynamical programming principle. $\qquad\square$

APPENDIX C. IMPLEMENTATION DETAILS

We are going to discuss the technical steps which are necessary in order to allow Monte-Carlo simulation for the signature stopping methodology. Throughout this section we will assume that the rough path $\mathbb{X} \in \Omega_T$ is the lift of a d-dimensional process $X = (X_t)_{0 \leq t \leq T}$ and that $Y = (Y_t)_{0 \leq t \leq T}$ is a continuous real-valued process adapted to the filtration generated by X.

C.1. **Time discretization.** We fix a time grid with equidistant points $(t_j)_{j=0,1,\ldots,J}$, $t_j = Tj/J$ and define the discretized payoff process $\widetilde{Y}$ by

$$\widetilde{Y}_j = Y_{t_j}, \quad j = 0, 1, \ldots, J.$$

Next, we fix a refinement $(s_j)_{j=0,1,\ldots,J'}$ of the grid $(t_j)_{j=0,1,\ldots,J}$ with $J' \geq J$ and denote by $\mathcal{S}_J$ the set of discrete stopping times with respect to the following filtration

$$\widetilde{\mathcal{F}}_j := \sigma\left(X_{s_i} \,:\, s_i \leq t_j, \, i = 0, 1, \ldots J'\right), \quad j = 0, 1, \ldots, J.$$

We can approximate the value of time optimal stopping problem (1.1) from below by the value of discrete time optimal stopping problem associated to $(\widetilde{Y}, (\widetilde{\mathcal{F}}_j)_{j=0,1,\ldots,J})$, i.e.

$$\sup_{\sigma \in \mathcal{S}_J} \mathbb{E}\left[\widetilde{Y}_{\sigma \wedge J}\right] \leq \sup_{\tau \in \mathcal{S}} \mathbb{E}\left[Y_{\tau \wedge T}\right].$$

Indeed, since $\widetilde{\mathcal{F}}_j \subset \mathcal{F}_{t_j}$ for all $j \in \{0, 1, \ldots, J\}$, it is easy to see that for any $\sigma \in \mathcal{S}_J$, $\omega \mapsto t_{\sigma(\omega)}$ defines a stopping time in $\mathcal{S}$. Next, denote by $\widetilde{X}$ the linear interpolation of the process X on the grid $(s_j)_{j=0,1,\ldots,J'}$, i.e.

$$\widetilde{X}_t = X_{s_j} + \frac{t - s_j}{s_{j+1} - s_j}(X_{s_{j+1}} - X_{s_j}), \quad s_j \leq t \leq s_{j+1}, \, j = 0, 1, \ldots, J'.$$

Clearly also the time augmented path $((t, \widetilde{X}_t))_{0 \leq t \leq T}$ is piecewise linear and this process has a (trivial) lift to a rough path, which we denote by $\widetilde{\mathbb{X}}$. We will now adapt the definition of the randomized stopping times (4.1) to the discrete time setting. Therefore, let Z be positive random variable as in Definition 4.1 and let $\theta \in \mathcal{T}$ be a continuous stopping rule, then we define

$$(C.1) \qquad \sigma_\theta^r = \inf\left\{ 0 \leq j \leq J \,\middle|\, \sum_{i=0}^{j} \theta(\widetilde{\mathbb{X}}|_{[0,t_i]})^2 \geq Z \right\}.$$

Since $\widetilde{\mathbb{X}}|_{[0,t_j]}$ is $\widetilde{\mathcal{F}}_j$-measurable for all $j \in \{0, 1, \ldots, J\}$, it follows that that $\sigma_\theta^r \in \mathcal{S}_J$. With minor changes in proof, an analogous result of Proposition 4.4 holds for the stopping times σ_θ^r and yields the following expression for the expected payoff

$$(C.2) \quad \mathbb{E}\left[\widetilde{Y}_{\sigma_\theta^r}\right] = \mathbb{E}\left[\widetilde{Y}_0 + \sum_{j=0}^{J-1} G(j)(\widetilde{Y}_{j+1} - \widetilde{Y}_j)\right], \qquad \text{with } G(j) := 1 - F_Z\left(\sum_{i=0}^{j} \theta(\widetilde{\mathbb{X}}|_{[0,t_i]})^2\right),$$

where F_Z is the cumulative distribution function of Z.

C.2. **Simulation of the (log-)signature.** The standard, model independent procedure for approximating the truncated signature of given simulations of a path is to calculate the truncated signature of the linearly interpolated path. Regarding the optimal stopping problem, it is therefore necessary to generate joint samples of the process X and the payoff Y, while one may possibly choose a finer approximation grid for the simulation of X. This may require discretization of stochastic / rough differential equations. The truncated signature of the linearly interpolated path on the given time grid can then be calculated exactly (up to floating point errors) and therefore always yields an element in the free nilpotent Lie-group. Fortunately, packages for this task are readily available. See, for instance, the *iisignature* library [RG20]. Note that in case the signature is needed on each subinterval $[0, s_j]$ (see Section C.4 below), one can use the calculation of the signature on the preceding subinterval to significantly improve the efficiency. As already mentioned in Section 7, the log-signature of the linearly interpolated path can be directly calculated in the Hall basis and does not require the calculation of the signature.

C.3. **Optimization of linear signature stopping policies based on the expected signature.** If the setting allows to pose the optimization problem in terms of the expected signature only, as in the context of Corollary 6.6, then we need to compute the expected truncated signature $\mathbb{E}[\widehat{\mathbb{X}}_{0,T}^{\leq N}]$. This can be done using Monte-Carlo estimation and the procedure for calculating samples of the truncated signature as in explaind in Section C.2 above. The big advantage is, of course, that we only need to compute the expected signature once, and can then apply the optimization algorithm of our choice to a deterministic optimization problem. For calculating the expression $\exp^{\sqcup\!\sqcup}(-(l \sqcup\!\sqcup l)1)$ upto truncation degree N, one needs to implement a structure that allows the handling of noncommutative polynomials with variable coefficients and that supports the shuffle operation. Since we could not find an appropriate package we have implemented this part ourself. Now let $\lambda_0 w_0 + \lambda_1 w_1 + \cdots + \lambda_{n_k} w_{n_k} \in T^N((\mathbb{R}^{d+1})^*)$, where $w_0, \ldots, w_{n_k}$ ranges over all words in of length $k \leq N$ in the alphabet $\mathcal{A}_{d+1}$ (a total number of $n_k + 1 = ((d+1)^{k+1} - 1)/d)$ words). Further, let $E \in T^N(\mathbb{R}^{d+1})$ be a Monte-Carlo estimate of the expected truncated signature signature, then the expression that needs to maximized is given by

$$\left\langle \exp^{\sqcup\!\sqcup} \left(-(\lambda_0 w_0 + \cdots + \lambda_{n_k} w_{n_k})^{\sqcup\!\sqcup 2} 1 \right) l_Y, E \right\rangle,$$

where $l_Y \in T((\mathbb{R}^{d+1})^*)$, such that $Y = \langle l_Y, \widehat{\mathbb{X}}^{<\infty} \rangle$. This expression is a polynomial in the coefficients $\lambda_1, \ldots, \lambda_{n_k}$. We can then apply a general state-of-the-art iterative optimization algorithms to optimize the above polynomial subject to the constraint $|\lambda_1| + \cdots + |\lambda_{n_k}| \leq K$. The resulting optimal coefficients can further be used to calculate a low-biased estimate of the optimal stopping problem as described in Section C.5 below.

C.4. **A loss-function for linear and deep signature stopping policies.** In this subsection we will define a loss-function for the linear and deep stopping policies based on the regular form of the expected payoff (C.2). The loss-function can be used with a stochastic gradient descent method in order to numerically optimize the parameters of the stopping policy and thus to approximate the value of the optimal stopping problem. Assume that we have generated M sample trajectories

$$(\widetilde{X}_j^{(m)}, \widetilde{Y}_j^{(m)})_{0 \leq j \leq J, \ 1 \leq m \leq M}, \qquad \text{with } \widetilde{X}_j^{(m)} \in \mathbb{R}^{d+1}, \quad \widetilde{Y}_j^{(m)} \in \mathbb{R},$$

of the time augmented path $\widehat{X}$ and the payoff process Y (for simplicity we have not used a finer time grid for $\widehat{X}$). Fix a truncation level $N \in \{1, 2, \ldots\}$. For each series of samples $\widetilde{X}_0^{(m)}, \ldots, \widetilde{X}_J^{(m)}$ we

can calculated the truncated signature of the linearly interpolated path on each subinterval $[0, t_j]$, yielding the the sequence of vectors

$$(\widetilde{\mathbb{X}}_j^{(m)})_{0 \le j \le J, \, 1 \le m \le M}, \qquad \text{with } \widetilde{\mathbb{X}}_j^{(m)} \in T^N(\mathbb{R}^{d+1}) \cong \mathbb{R}^{1+\sigma_{d+1,N}}.$$

Let $\theta_l \in \mathcal{T}_{\text{sig}}$ be a signature stopping policy of the form $\theta_l(\widehat{\mathbb{X}}) = \langle l, \widehat{\mathbb{X}}^{<\infty} \rangle$ for some $l \in T^N((\mathbb{R}^{d+1})^*)$. Further let $\mathcal{M} \subset \{1, ..., M\}$ be the indices of a batch of samples. Using (C.2) we define the following loss-function for the stopping policy θ_l

$$\ell\left(\widetilde{Y}|_{\mathcal{M}}, \widetilde{\mathbb{L}}|_{\mathcal{M}}; \, l\right) = -\frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} \left\{ \widetilde{Y}_0^{(m)} + \sum_{j=0}^{J-1} G_Z \left(\sum_{i=0}^{j} \left\langle l, \widetilde{\mathbb{X}}_i^{(m)} \right\rangle^2 \right) (\widetilde{Y}_{j+1}^{(m)} - \widetilde{Y}_j^{(m)}) \right\},$$

where $G_Z = 1 - F_Z$.

Next, let $\theta = \theta_{\log} \circ \log^{\otimes} \in \mathcal{T}_{\log}$ be a deep stopping policy with a neural network $\theta_{\log}$ of the form (7.2). In this case, the definition of the loss-function is in principle the same as above, however, the log-signature can be calculated directly in the hall basis, yielding the sequence of vectors

$$(\widetilde{\mathbb{L}}_j^{(m)})_{0 \le j \le J, \, 1 \le m \le M}, \qquad \text{with } \widetilde{\mathbb{L}}_j^{(m)} \in \mathfrak{g}^N(\mathbb{R}^{d+1}) \cong \mathbb{R}^{\eta_{d+1,N}}.$$

Again using (C.2) we define the following loss-function for the deep stopping policy θ by

$$\ell\left(\widetilde{Y}|_{\mathcal{M}}, \widetilde{\mathbb{L}}|_{\mathcal{M}}; \theta_{\log}\right) = -\frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} \left\{ \widetilde{Y}_0^{(m)} + \sum_{j=0}^{J-1} G_Z \left(\sum_{i=0}^{j} \theta_{\log} \left(\widetilde{\mathbb{L}}_i^{(m)} \right)^2 \right) (\widetilde{Y}_{j+1}^{(m)} - \widetilde{Y}_j^{(m)}) \right\}.$$

Regarding the distribution of Z, we have tried the exponential distribution and the log-logistic distribution. The latter, corresponding to $G_Z(x) = 1/(1+x)$, led to initially quicker learning rates, however for the overall performance it was better to use the exponential distribution.

C.5. **Calculation of lower-bounds.** Assume that we have obtained a stopping policy $\theta^* \in \mathcal{T}$ as the output of a numerical optimization. For example $\theta^* = \langle l^*, \cdot \rangle \in \mathcal{T}_{\text{sig}}$ for some numerically optimized $l^* \in T(V^*)$; or $\theta^* \in \mathcal{T}_{\log}$ with some trained neural network $\theta_{\log}^*$. We then calculate a estimate to the value of the optimal stopping problem by a Monte-Carlo approximation of (C.2) with a new set of sample trajectories. Note that this estimates is precisely given by the respective loss-function of Section C.4 evaluated at the new set of samples. The estimate is low biased due to the sub-optimality of the stopping policy.

Alternatively, note that the stopping policy θ^* defines a stopping time $\sigma^* := \sigma_{\theta^*}^r \in \mathcal{S}_J$ by (C.1). Therefore, after generating independent samples of Z, one for each sample trajectory of X and Y,

we obtain another low-biased estimator of the optimal stopping value by the Monte-Carlo approximation of $\mathbb{E}[\widetilde{Y}_{\sigma^*}]$. Since both estimates yield the same result, upto a the additional Monte-Carlo error introduced by the sampling of Z, the latter value can serve for a sanity check.

Christian Bayer, Weierstrass Institute, Berlin, Germany

Email address: christian.bayer@wias-berlin.de

Paul Hager, Institut für Mathematik, Technische Universität Berlin, Germany

Email address: phager@math.tu-berlin.de

Sebastian Riedel, Institut für Analysis, Leibniz Universität Hannover, Germany

Email address: riedel@math.uni-hannover.de

John Schoenmakers, Weierstrass Institute, Berlin, Germany

Email address: john.schoenmakers@wias-berlin.de