Signatures in machine learning and finance

Imanol Perez Arribas
Keble College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2020

Acknowledgements

First and foremost, I would like to thank my supervisor Terry Lyons for his guidance throughout my DPhil. Our conversations (both mathematical and otherwise) have shaped me into much of what I am today. I would also like to thank my colleagues at J.P. Morgan for the opportunity I was given to learn more about mathematical finance in an industrial environment. I am grateful to my collaborators and colleagues for the many talks we have had over these past years: they have helped me reach further than I ever would on my own.

This work would not have been possible without the support of EPSRC, The Alan Turing Institute and La Caixa.

Last but not least, I would like to thank the infinite support of my family and Marina: they have been part of this journey as much as I have, for which I will be forever grateful.

Abstract

Developed to give meaning to differential equations driven by rough signals, rough path theory has opened in recent years a new approach to tackle certain problems in other fields such as mathematical finance and machine learning. This is due to certain algebraic and analytical properties of an object called the *rough path signature*. This thesis aims to make a contribution in this direction by introducing signature-based methods to study various problems arising in finance and machine learning.

In the context of finance, we consider two problems: the pricing and hedging of financial derivatives, and the optimal execution of financial securities. We propose numerical methods based on signatures to solve both problems.

Finally, in the context of machine learning we begin by introducing a new neural network architecture called the *deep signature model*, based on the *deep signature transform*. The thesis concludes with a practical application of the signature transform on a real-world problem, where we aim to make predictions from psychiatric data using signatures.

List of symbols and notation

V	a finite-dimensional vector space
V^*	dual of V
$T((V))$	extended tensor algebra over V
$T(V)$	tensor algebra over V
$T^{(n)}(V)$	tensor algebra of order $n \in \mathbb{N}$ over V
π_n	canonical projection $T((V)) \to V^{\otimes n}$
Π_n	canonical projection $T((V)) \to T^{(n)}(V)$
$\mathbb{X}^N$	N-th term of the signature
$\mathbb{X}^{\leq N}$	signature up to level N
$\mathbb{X}^{<\infty}$	full signature
$\sqcup\!\sqcup$	shuffle product
$\ell^{\sqcup\!\sqcup n}$	notation for $\underbrace{\ell \sqcup\!\sqcup \ldots \sqcup\!\sqcup \ell}_{n}$
$\mathcal{A}_d$	alphabet of $d \in \mathbb{N}$ letters
$\mathcal{W}(\mathcal{A}_d)$	the vector space of all words with letters in the alphabet $\mathcal{A}_d$
$\mathbf{w}$	a generic word
$\varnothing$	the empty word
$G^{(*)}(V)$	space of grouplike elements
$G^{(n)}(V)$	projection of $G^{(*)}(V)$ to $T^{(n)}(V)$
d_{p-var}	p-variation distance
$G\Omega_p([0,T];V)$	space of geometric p-rough paths
$(\mathbb{Z}^{\leq 2}, \langle Z \rangle)$	lead-lag path
$\mathbb{Z}^{LL,<\infty}$	signature of the lead-lag path $(\mathbb{Z}^{\leq 2}, \langle Z \rangle)$
$\mathcal{T}$	space of trading strategies
$\mathcal{T}_{sig}$	space of linear signature trading strategies
$\mathcal{S}(\mathcal{X})$	space of streams of data in a set $\mathcal{X}$
$\mathbf{x}$	a generic element $\mathbf{x} \in \mathcal{S}(\mathcal{X})$
Sig	signature layer
Sig^N	truncated signature layer of order $N \in \mathbb{N}$

Contents

Preface

The main objective of my research is to explore different applications of rough path signatures. Developed in the 1990s by Terry Lyons, rough path theory aimed to study differential equations driven by rough signals. However, over the past decade the use of signatures in various applications such as data science has opened a new line of research.

My goal is to apply signatures in two fields that have gained popularity lately: mathematical finance and machine learning. The manuscript is divided into three parts. The first part introduces signatures and rough path theory. Although a full introduction to the theory of rough paths is beyond the scope of this manuscript, a formal definition of what a rough path is and what some of its properties are is necessary. The reader is referred to [105, 56] for a more detailed review of rough path theory.

The second part of this manuscript discusses some recent work regarding applications of signatures in finance. It was written around the papers [101, 102, 86, 27]. The first two introduce a family of path-dependent financial derivatives – called *signature payoffs* – that are used to price and hedge in a model-free manner, either in continuous or discrete time, other exotic payoffs. This is explored in Chapters 3 and 4. The third paper [86], on which Chapter 5 is based, proposes a numerical method based on signatures to solve certain optimal execution problems that arise in algorithmic trading, and [27] builds on these results to solve a certain *double-execution* problem

where the investor aims to liquidate a block of shares in a foreign stock market and the gains from the sale are exchanged into the investor's domestic currency. To the best of our knowledge, this is the first time this problem is studied and solved in the literature. The last chapter of Part II, Chapter 6, is devoted to this double-execution problem.

Finally, the third part aims to explore applications of signatures in machine learning. The first chapter of Part III, based on the paper [87], proposes a new neural network architecture for streamed data, called the *deep signature transform*. The second chapter of the third part of the manuscript shows a real-world application of signatures. The chapter, based on the paper [117], applies signatures to a machine learning problem from the field of mental health.

Part I

Signatures

Chapter 1

Tensor algebras and their dual

1.1 Tensor algebra

Tensor algebras are formal series of tensors that are equipped with certain algebraic operations. A precise definition is given below. Hereafter V will denote a d-dimensional real vector space.

Definition 1.1.1 (Extended tensor algebra)**.** *We define the extended tensor algebra over V by*

$$T((V)) := \prod_{k=0}^{\infty} V^{\otimes k}.$$

We equip $T((V))$ with an addition $+$ and a product $\otimes$. For $\mathbf{a}, \mathbf{b} \in T((V))$ given by $\mathbf{a} = (a_n)_{n=0}^{\infty}$, $\mathbf{b} = (b_n)_{n=0}^{\infty}$ with $a_n, b_n \in V^{\otimes n}$ for each $n \geq 0$, these two operations are defined by

$$\mathbf{a} + \mathbf{b} := (a_0 + b_0, a_1 + b_1, \ldots) \in T((V))$$

and

$$\mathbf{a} \otimes \mathbf{b} := (c_0, c_1, \ldots) \in T((V))$$

where

$$c_n := \sum_{k=0}^{n} a_k \otimes b_{n-k} \in V^{\otimes n} \quad \forall n \geq 0.$$

We also define the action on $\mathbb{R}$ given by $\lambda \mathbf{a} := (\lambda a_n)_{n=0}^{\infty}$ for all $\lambda \in \mathbb{R}$.

As the name suggests, $T((V))$, equipped with the sum $+$ and product $\otimes$, is an algebra with unit $\mathbf{1} := (1, 0, 0, \ldots)$. Notice that unless V is one-dimensional, $T((V))$ is non-commutative.

Intuitively the extended tensor algebra consists of all sequences of tensors. Similarly, we may define the subalgebra of all *finite* sequences and all sequences *of a given length*, which will be called the tensor algebra and truncated tensor algebra respectively.

Definition 1.1.2. *The tensor algebra over V, denoted by $T(V) \subset T((V))$, is given by*

$$T(V) := \bigoplus_{k=0}^{\infty} V^{\otimes k}.$$

Similarly, the truncated tensor algebra of order $n \in \mathbb{N}$ over V is defined by

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

For each $n \geq 1$, we define the projection

$$\begin{aligned} \pi_n : T((V)) &\to V^{\otimes n} \\ (a_i)_{i=0}^{\infty} &\mapsto a_n \in V^{\otimes n}. \end{aligned}$$

We also denote the canonical projection $T((V)) \to T^{(n)}(V)$ by Π_n:

$$\begin{aligned} \Pi_n : T((V)) &\to T^{(n)}(V) \\ (a_i)_{i=0}^{\infty} &\mapsto (a_i)_{i=0}^{n} \in T^{(n)}(V). \end{aligned}$$

Remark 1.1.3. If V is a vector space of dimension $d > 1$, then $T^{(n)}(V)$ has dimension $1 + d + d^2 + \ldots + d^n = (d^{n+1} - 1)/(d - 1)$.

1.2 The dual of the tensor algebra

Let $\{e_1, \ldots, e_d\} \subset V$ be a basis for V. This induces a dual basis $\{e_1^*, \ldots, e_d^*\} \subset V^*$, where V^* denotes the dual space of V. We may define a basis for $V^{\otimes n}$ by:

$$\{e_{i_1} \otimes \ldots \otimes e_{i_n} \mid i_j \in \{1, \ldots, d\} \text{ for } j = 1, \ldots, n\}.$$

Similarly, a basis for $(V^*)^{\otimes n}$ is defined by

$$\{e_{i_1}^* \otimes \ldots \otimes e_{i_n}^* \mid i_j \in \{1, \ldots, d\} \text{ for } j = 1, \ldots, n\}.$$

This induces, in a natural way, a basis for $T((V))$ and $T(V^*)$. In fact, we have $T((V))^* \cong T(V^*)$ (see [105]).

1.3 Words

It is often convenient to think of $T(V^*)$ as a space of *words*. Define the alphabet $\mathcal{A}_d := \{\mathbf{1}, \ldots, \mathbf{d}\}$. Then, the basic element $e_{i_1}^* \otimes \ldots \otimes e_{i_n}^* \in T(V^*)$ can be identified with the word $\mathbf{i_1} \ldots \mathbf{i_n}$. Let $\mathcal{W}(\mathcal{A}_d)$ denote the vector space of all words with letters in the alphabet $\mathcal{A}_d$, i.e. the $\mathbb{R}$-vector space generated by $\mathcal{A}_d$. Then, we have $T(V^*) \cong \mathcal{W}(\mathcal{A}_d)$. The empty word will be denoted by $\varnothing \in \mathcal{W}(\mathcal{A}_d)$.

Example 1.3.1. Consider the following examples for $V = \mathbb{R}^2$.

1. Let $\mathbf{a} = 3 - e_2 \otimes e_1 \in T((\mathbb{R}^2))$. Then, $\langle \varnothing, \mathbf{a} \rangle = 3$.
2. Let $\mathbf{a} = 1 - 2e_1 + e_2 \in T((\mathbb{R}^2))$, and set $\mathbf{w} = \varnothing + \mathbf{1}$. Then, $\langle \mathbf{w}, \mathbf{a} \rangle = 1 - 2 = -1$.
3. Let $\mathbf{a} = e_1 \otimes e_2 - e_2 \otimes e_1 \in T((\mathbb{R}^2))$, and set $\mathbf{w} = \mathbf{21} + \mathbf{111}$. Then, $\langle \mathbf{w}, \mathbf{a} \rangle = -1 + 0 = -1$.
4. Let $\mathbf{a} = 1 + e_1^{\otimes 3}$ and $\mathbf{w} = 2 \cdot \mathbf{111}$. Then, $\langle \mathbf{w}, \mathbf{a} \rangle = 2 \cdot 1 = 2$.

The space of words possesses two natural algebraic operations – the sum and the concatenation. Let $\mathbf{w} = \mathbf{i_1} \dots \mathbf{i_n}, \mathbf{v} = \mathbf{j_1} \dots \mathbf{j_m} \in \mathcal{W}(\mathcal{A}_d)$ be two words. Their sum is just the formal sum $\mathbf{w} + \mathbf{v} \in \mathcal{W}(\mathcal{A}_d)$. Their concatenation, on the other hand, is defined by

$$\mathbf{wv} := \mathbf{i_1} \dots \mathbf{i_n j_1} \dots \mathbf{j_m} \in \mathcal{W}(\mathcal{A}_d)$$

which is extended by bilinearity to $\mathcal{W}(\mathcal{A}_d)$. These two operations induce analogous operations on $T(V^*)$, and with some abuse of notation we will even use concatenation on $\mathcal{W}(\mathcal{A}_d)$ and $T(V^*)$ interchangeably – i.e. we will sometimes write $\ell\mathbf{w} \in T(V^*)$ for $\ell \in T(V^*)$ and word $\mathbf{w} \in \mathcal{W}(\mathcal{A}_d)$, by which we mean that we take the concatenation of the element in $\mathcal{W}(\mathcal{A}_d)$ associated to ℓ and the word $\mathbf{w}$. For example, if $\ell = e_2^* \in T((\mathbb{R}^2)^*)$ and $\mathbf{w} = \mathbf{12} \in \mathcal{W}(\mathcal{A}_2)$, then $\ell\mathbf{w}$ is just $e_2^* \otimes e_1^* \otimes e_2^* \in T((\mathbb{R}^2)^*)$, or using the notation of words, it is just $\mathbf{212} \in \mathcal{W}(\mathcal{A}_2)$.

Example 1.3.2. Take the alphabet $\mathcal{A}_4 = \{\mathbf{1}, \mathbf{2}, \mathbf{3}, \mathbf{4}\}$.

1. Set $\mathbf{w} = \mathbf{212}, \mathbf{v} = \mathbf{31}$. We have $\mathbf{wv} = \mathbf{21231} \in \mathcal{W}(\mathcal{A}_4)$.

2. We have $(\mathbf{143} + \mathbf{23})\mathbf{1} = \mathbf{1431} + \mathbf{231} \in \mathcal{W}(\mathcal{A}_4)$.

There is a third operation on words that will be useful in this thesis: the *shuffle product* ш . Intuitively, the shuffle product accounts for all the possible ways of riffle shuffling two decks of cards. The precise definition is given below.

Definition 1.3.3 (Shuffle product)**.** *The shuffle product* $\text{ш} : \mathcal{W}(\mathcal{A}_d) \times \mathcal{W}(\mathcal{A}_d) \to \mathcal{W}(\mathcal{A}_d)$ *is defined inductively by*

$$\mathbf{ua} \text{ш} \mathbf{vb} = (\mathbf{u} \text{ш} \mathbf{vb})\mathbf{a} + (\mathbf{ua} \text{ш} \mathbf{v})\mathbf{b},$$

$$\mathbf{w} \text{ш} \varnothing = \varnothing \text{ш} \mathbf{w} = \mathbf{w}$$

for all words $\mathbf{u}, \mathbf{v}, \mathbf{w}$ *and letters* $\mathbf{a}, \mathbf{b} \in \mathcal{A}_d$*, which is then extended by bilinearity to* $\mathcal{W}(\mathcal{A}_d)$*. With some abuse of notation, the shuffle product on* $T(V^*)$ *induced by the shuffle product on words will also be denoted by* $\sqcup\!\sqcup$*.*

It follows from the definition of the shuffle product that $\sqcup\!\sqcup$ is commutative, i.e. $f \sqcup\!\sqcup g = g \sqcup\!\sqcup f$ for all $f, g \in T(V^*)$.

Example 1.3.4. We have:

1. $\mathbf{12} \sqcup\!\sqcup \mathbf{3} = \mathbf{123} + \mathbf{132} + \mathbf{312}$.

2. $\mathbf{12} \sqcup\!\sqcup \mathbf{24} = 2 \cdot \mathbf{1224} + \mathbf{1242} + \mathbf{2124} + \mathbf{2142} + \mathbf{2412}$.

Definition 1.3.5. *Let* $Q \in \mathbb{R}[x]$ *be a polynomial on one variable. Write* $Q(x) = a_0 + a_1 x + a_2 x^2 + \ldots + a_n x^n$*. Then,* Q *induces the map* $Q^{\sqcup\!\sqcup} : T(V^*) \to T(V^*)$ *given by*

$$Q^{\sqcup\!\sqcup}(\ell) := a_0 \varnothing + a_1 \ell + a_2 \ell^{\sqcup\!\sqcup 2} + \ldots + a_n \ell^{\sqcup\!\sqcup n} \quad \forall \ell \in T(V^*),$$

where $\ell^{\sqcup\!\sqcup k} := \underbrace{\ell \sqcup\!\sqcup \ell \sqcup\!\sqcup \ldots \sqcup\!\sqcup \ell}_{k}$ *for* $k \in \mathbb{N}$.

Chapter 2

Rough path theory

2.1 Signatures of smooth paths

Throughout this chapter, V will denote a finite-dimensional Banach space with norm $\|\cdot\|$. In Chapter 1 we introduced the *tensor algebra* $T((V))$, a graded space of tensors equipped with a certain sum and a certain product that give it an algebra structure. The signature of a (smooth) path $X : [0,T] \to V$ is an element in this space. A precise definition is given below.

Definition 2.1.1 (Signature of a smooth path)**.** *Let $0 \leq s < t \leq T$. For a piecewise smooth path $X : [0,T] \to V$, we define the signature of X over $[s,t]$ by*

$$\mathbb{X}_{s,t}^{<\infty} := (1, \mathbb{X}_{s,t}^1, \dots, \mathbb{X}_{s,t}^n, \dots) \in T((V))$$

where

$$\mathbb{X}_{s,t}^n := \int_{s<u_1<\ldots<u_n<t} dX_{u_1} \otimes \ldots \otimes dX_{u_n} \in V^{\otimes n}. \tag{2.1}$$

Similarly, we define the truncated signature of order $N \in \mathbb{N}$ by

$$\mathbb{X}_{s,t}^{\leq N} := (1, \mathbb{X}_{s,t}^1, \dots, \mathbb{X}_{s,t}^N) \in T^{(N)}(V).$$

If we refer to the signature of X, without referencing the interval over which the signature is taken, we will implicitly refer to $\mathbb{X}_{0,T}^{<\infty}$.

Example 2.1.2. Linear functions on signatures will play a key role in later chapters. Therefore, it will be useful to see a few examples that will be used later.

Let $X = (X^1, X^2) \in C^\infty([0,T];\mathbb{R}^2)$ be a two-dimensional smooth path. Recall that in Section 1.3 we introduced the notation of *words* as linear functions on the tensor algebra. We have:

1. $\langle \mathbf{2}, \mathbb{X}_{0,T}^{<\infty} \rangle = \int_0^T dX_t^2 = X_T^2 - X_0^2$.
2. $\langle \varnothing, \mathbb{X}_{0,T}^{<\infty} \rangle = 1$.
3. $\langle \mathbf{21}, \mathbb{X}_{0,T}^{<\infty} \rangle = \int_0^T \int_0^t dX_s^2 dX_t^1 = \int_0^T (X_t^2 - X_0^2) dX_t^1$.
4. Let $\ell \in T((\mathbb{R}^2)^*)$. Then, $\langle \ell \mathbf{1}, \mathbb{X}_{0,T}^{<\infty} \rangle = \int_0^T \langle \ell, \mathbb{X}_{0,t}^{<\infty} \rangle dX_t^1$.

We will now state below two important properties of signatures.

Proposition 2.1.3 (Invariance under reparametrisation, [105])**.** *Let $X : [0,T] \to V$ be a piecewise smooth path, and let $\varphi : [0,T] \to [0,T]$ be a reparametrisation. Then, the signature of X coincides with the signature of $X \circ \varphi$.*

Proposition 2.1.4 (Chen's identity, [105])**.** *Let $X, Y : [0,T] \to V$ be two piecewise smooth paths and denote by $X * Y$ their concatenation, which is given by the (piecewise smooth) path $X * Y : [0,T] \to V$ defined by*

$$(X * Y)_t := \begin{cases} X_{2t}, & t \in [0, T/2]; \\ Y_{2t-T} - Y_0 + X_T, & t \in [T/2, T]. \end{cases}$$

Then, the signature of $X * Y$ is given by $\mathbb{X}_{0,T}^{<\infty} \otimes \mathbb{Y}_{0,T}^{<\infty}$.

Remark 2.1.5. Chen's identity makes the computation of signatures of discrete data (say, piecewise linear paths) efficient. By Chen's identity the signature of a piecewise linear path is given by the tensor product of the signatures of each linear piece, and because signatures of linear paths can be computed explicitly (see [95]) computing the signature of the whole piecewise linear path is computationally cheap.

Let $X : [0,T] \to V$ be a smooth path. By definition, its signature $\mathbb{X}_{0,T}^{<\infty}$ is an element of $T((V))$. However, it is not *a priori* clear in what subspace of $T((V))$ signatures of smooth paths live. It turns out that signatures take values in a certain group called the space of *grouplike elements*.

Definition 2.1.6 (Grouplike elements)**.** *An element $\mathbf{a} \in T((V)) \setminus \{0\}$ is said to be grouplike if given any $\ell_1, \ell_2 \in T(V^*)$, the following equality holds:*

$$\langle \ell_1, \mathbf{a} \rangle \langle \ell_2, \mathbf{a} \rangle = \langle \ell_1 \text{ш} \ell_2, \mathbf{a} \rangle$$

where ш denotes the shuffle product (Definition 1.3.3).
The space of all grouplike elements $\mathbf{a} \in T((V)) \setminus \{0\}$ is denoted by $G^{(*)}(V)$. We also define, for $n \in \mathbb{N}$, the projection of $G^{(*)}(V)$ into $T^{(n)}(V)$:

$$G^{(n)}(V) := \Pi_n \left(G^{(*)}(V) \right) \subset T^{(n)}(V).$$

Proposition 2.1.7 (Shuffle product property, [105])**.** *Let $X : [0,T] \to V$ be a smooth path and let $\mathbb{X}_{0,T}^{<\infty}$ be its signature. Then, $\mathbb{X}_{0,T}^{<\infty} \in G^{(*)}(V)$, i.e.*

$$\langle \ell_1, \mathbb{X}_{0,T}^{<\infty} \rangle \langle \ell_2, \mathbb{X}_{0,T}^{<\infty} \rangle = \langle \ell_1 \text{ш} \ell_2, \mathbb{X}_{0,T}^{<\infty} \rangle$$

for all $\ell_1, \ell_2 \in T(V^*)$.

By the shuffle product property of signatures, polynomials on signatures turn out to be linear functions on signatures.

2.2 Geometric rough paths

Definition 2.2.1 (p-variation distance)**.** *For $T > 0$, denote by $\Delta_T := \{(s,t) \in [0,T]^2 \mid s \leq t\}$. Set $p \geq 1$, and let $\lfloor p \rfloor$ be its integer part. Given $\mathbb{X}, \mathbb{Y} : \Delta_T \to G^{\lfloor p \rfloor}(V)$, their p-variation distance is defined by*

$$d_{p-var}(\mathbb{X}, \mathbb{Y})^p := \sup_D \sum_{t_i \in D} \sum_{k=1}^{\lfloor p \rfloor} \left\| \pi_k \left(\left(\mathbb{X}_{t_{i-1},t_i} \right)^{-1} \mathbb{Y}_{t_{i-1},t_i} \right) \right\|_{V^{\otimes k}}^{p/k}$$

where the supremum is taken over all finite partitions D of $[0,T]$.

Definition 2.2.2 (Geometric p-rough paths)**.** *Let $T > 0$ and $p \geq 1$. The space of geometric p-rough paths $G\Omega_p([0,T];V)$ is defined as the closure (under the p-variation distance) of signatures of order $\lfloor p \rfloor$ of smooth paths.*

Theorem 2.2.3. *[Extension theorem, [105, Theorem 3.7]] Let $\mathbb{X} \in G\Omega_p([0,T];V)$. Given any $N \in \mathbb{N}$ with $N > \lfloor p \rfloor$, there exists a unique geometric N-rough path $\mathbb{X}^{\leq N} \in G\Omega_N([0,T];V)$ such that*

(i) $\mathbb{X}^{\leq N}$ is an extension of $\mathbb{X}$, i.e. $\Pi_{\lfloor p \rfloor}(\mathbb{X}^{\leq N}) = \mathbb{X}$.

(ii) $\mathbb{X}^{\leq N}$ has finite p-variation, i.e.

$$\sup_D \sum_{t_i \in D} \sum_{k=1}^{N} \left\| \pi_k \left(\mathbb{X}^{\leq N}_{t_i, t_{i+1}} \right) \right\|^{p/k}_{V^{\otimes k}} < \infty$$

where the supremum is taken over all finite partitions D of $[0,T]$.

The extension theorem may be stated for more general rough paths (see [105, Theorem 3.7]]) but the statement above will suffice for our purpose.

Definition 2.2.4 (The signature)**.** *Let $p \geq 1$, and let $\mathbb{X} \in G\Omega_p([0,T];V)$ be a p-geometric rough path. Given $N \in \mathbb{N}$ such that $N > \lfloor p \rfloor$, we call the extension $\mathbb{X}^{\leq N} = (1, \mathbb{X}^1, \ldots, \mathbb{X}^N) : \Delta_T \to G^{(N)}(V) \subset T^{(N)}(V)$ from Theorem 2.2.3 the level-N truncated signature of $\mathbb{X}$.*

Likewise, we refer to $\mathbb{X}^{<\infty} = (1, \mathbb{X}^1, \ldots, \mathbb{X}^N, \ldots) : \Delta_T \to G^{(*)}(V) \subset T((V))$ as the signature of $\mathbb{X}$.

A geometric p-rough path $\mathbb{X} = (1, \mathbb{X}^1, \ldots, \mathbb{X}^{\lfloor p \rfloor}) \in G\Omega_p([0,T];V)$ is a continuous map $\mathbb{X} : \Delta_T \to G^{\lfloor p \rfloor}(V)$ that is the limit of truncated signatures of order $\lfloor p \rfloor$ of smooth paths. Its signature $\mathbb{X}^{<\infty}$ is grouplike, and hence it satisfies the shuffle product property

$$\langle \ell_1, \mathbb{X}^{<\infty}_{s,t} \rangle \langle \ell_2, \mathbb{X}^{<\infty}_{s,t} \rangle = \langle \ell_1 \text{ш} \ell_2, \mathbb{X}^{<\infty}_{s,t} \rangle \quad \forall\, (s,t) \in \Delta_T,\ \ell_1, \ell_2 \in T(V^*).$$

Example 2.2.5. Many stochastic processes that are used in the literature are almost surely geometric rough paths. For example, the signature of a semimartingale, defined via the iterated integrals in (2.1) and using Stratonovich integration, is almost surely a geometric p-rough path for any $p > 2$ [38]. The signature of a fractional Brownian motion for Hurst parameter $H \geq 1/4$, defined almost surely, is also a geometric p-rough path for $p > 1/H$ ([39]).

The following lemma and the subsequent corollary states one of the key properties of signatures. It guarantees that the signature $\mathbb{X}_{0,T}^{<\infty}$ completely characterises $\mathbb{X}$ – up to the so-called tree-like equivalences (see [15, Definition 1.1]).

Lemma 2.2.6 (Uniqueness of signatures, [15])**.** *Let $\mathbb{X} \in G\Omega_p([0,T];V)$. The signature $\mathbb{X}_{0,T}^{<\infty}$ of $\mathbb{X}$ is unique up to tree-like equivalences (defined in [15, Definition 1.1]).*

Corollary 2.2.7. *Let $\mathbb{X} \in G\Omega_p([0,T];V)$. If there exists a projection of $\mathbb{X}$ that is strictly monotone, then the signature $\mathbb{X}_{0,T}^{<\infty}$ determines $\mathbb{X}$.*

Example 2.2.8. Let $X : [0,T] \to V$ be smooth, and consider the path $\widehat{X} : [0,T] \to \mathbb{R} \oplus V$ given by $\widehat{X}_t := (t, X_t)$. Because time is strictly monotone, by the previous corollary the signature of $\widehat{X}$, i.e. $\widehat{\mathbb{X}}_{0,T}^{<\infty}$, uniquely determines X up to translations.

2.3 Lead-lag path

In this section we introduce *lead-lag paths*, also called Hoff processes. First studied by B. Hoff in [78] for Brownian paths and later by G. Flint, B. Hambly and T. Lyons in [52] for semimartingales, lead-lag paths offer a canonical way of transforming discrete streamed data into a rough path by introducing a certain volatility process.

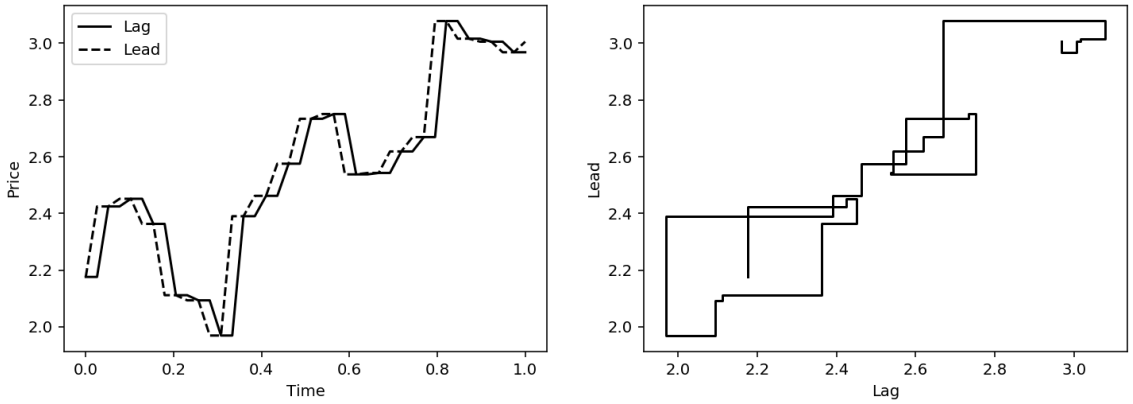

Figure 2.1: Lead-lag transformation of the price path process of a certain financial asset. The figure on the left shows the lead and lag components of the path, and the figure on the right shows the lag component plotted against the lead component.

Definition 2.3.1 (Lead-lag path)**.** *Let $p \in [2,3)$. A lead-lag path is a pair $(\mathbb{Z}^{\leq 2}, \langle Z \rangle)$ with $\mathbb{Z}^{\leq 2} = (1, \mathbb{Z}^1, \mathbb{Z}^2) \in G\Omega_p([0,T];V)$ a p-geometric rough path and $\langle Z \rangle : \Delta_T \to V \otimes V$ such that $\langle Z \rangle$ is symmetric and*

$$\mathbb{Z}^{LL,\leq 2} := \left(1, (\mathbb{Z}^1, \mathbb{Z}^1), \begin{pmatrix} \mathbb{Z}^2 & \mathbb{Z}^2 - \frac{1}{2}\langle Z \rangle \\ \mathbb{Z}^2 + \frac{1}{2}\langle Z \rangle & \mathbb{Z}^2 \end{pmatrix}\right) \in G\Omega_p([0,T]; V \oplus V)$$

is a geometric rough path on $V \oplus V$. $\mathbb{Z}^{LL,<\infty}$ will be called the signature of the lead-lag path $(\mathbb{Z}^{\leq 2}, \langle Z \rangle)$.

Example 2.3.2. A continuous semimartingale $Z : [0,T] \to \mathbb{R}^d$ induces a geometric rough path $\mathbb{Z}^{\leq 2}$, as shown in Example 2.2.5. This geometric rough path is given by certain Stratonovich iterated integrals. $\mathbb{Z}^{\leq 2}$, together with the quadratic variation $\langle Z \rangle_{s,t} \in \mathbb{R}^{d \times d}$ of Z over $[s,t]$, induces the lead-lag path $(\mathbb{Z}^{\leq 2}, \langle Z \rangle)$. Such lead-lag paths were considered in [52]. As we will see in Lemma 3.2.11, certain Itô integrals against the semimartingale can be written as integrals against the lead-lag process. This will prove useful in Chapter 3 to write the P&L of a trading strategy, given in terms of an Itô integral, as a function on the signature of the lead-lag path.

2.3.1 Computation of lead-lag paths

In this section, we will discuss some practical considerations about how to compute the signature of lead-lag paths for discrete data, as well as for semimartingales.

Let $D = \{t_i\}_{i=0}^n \subset [0,T]$ be a finite partition, and let $Z : D \to \mathbb{R}^d$ be a discrete path. The *lead-lag* transformation of Z is defined below.

Definition 2.3.3. *[Lead-lag transformation, [52, Definition 2.1]] The lead-lag transformation of Z associated with D is the $2d$-dimensional piecewise linear path $Z^{D,LL} := (Z^{D,b}, Z^{D,f}) : [0,T] \to \mathbb{R}^{2d}$ defined by*

$$Z_t^{D,LL} := \begin{cases} \left(Z_{t_k}, Z_{t_{k+1}}\right), & t \in \left[\frac{2k}{2n}T, \frac{2k+1}{2n}T\right), \\ \left(Z_{t_k}, Z_{t_{k+1}} + 2(t-(2k+1))\left(Z_{t_{k+2}} - Z_{t_{k+1}}\right)\right), & t \in \left[\frac{2k+1}{2n}T, \frac{2k+3/2}{2n}T\right), \\ \left(Z_{t_k} + 2(t-(2k+\frac{3}{2}))\left(Z_{t_{k+1}} - Z_{t_k}\right), Z_{t_{k+2}}\right), & t \in \left[\frac{2k+3/2}{2n}T, \frac{2k+2}{2n}T\right). \end{cases}$$

The component $Z^{D,b}$ is the lag or backward component, and $Z^{D,f}$ is the lead or forward component. By taking the signature of this piecewise linear path, we obtain the signature of the lead-lag path $\mathbb{Z}^{D,LL,<\infty}$.

The lead-lag transformation differentiates the role played by the *past* and the *future*. This is done by keeping track of the immediate past (the *lag* component) and the immediate future (the *lead* component).

In order to give an intuition of what the lead-lag transformation is, Figure 2.1 shows the lead-lag transformation of a certain price path of a financial asset. As the name suggests, the lead component is *leading* the lag component.

Now, let $Z : [0,T] \to \mathbb{R}^d$ be a continuous semimartingale with quadratic variation $\langle Z \rangle$. As discussed in Example 2.3.2, this induces a lead-lag path $(\mathbb{Z}^{\leq 2}, \langle Z \rangle)$ whose signature is $\mathbb{Z}^{LL,<\infty}$. It turns out (see the lemma below) that if we sample the semimartingale and compute the signature of the lead-lag transformation of the sampled path, this object converges to $\mathbb{Z}^{LL,<\infty}$ as the partition size goes to zero. This offers

a method to approximate the signature of the lead-lag path of a semimartingale in practice.

Lemma 2.3.4. *[[52, Theorem 4.1]] Let $Z : [0,T] \to \mathbb{R}^d$ be a continuous semimartingale. For each finite partition $D \subset [0,T]$, denote by $Z^{D,LL}$ the lead-lag transformation of $\{Z_t\}_{t\in D}$ and by $\mathbb{Z}^{D,LL,<\infty}$ its signature (Definition 2.3.3). Let $\mathbb{Z}^{LL,<\infty}$ be the signature of the lead-lag path $(\mathbb{Z}^{\leq 2}, \langle Z\rangle)$ associated with the semimartingale (Definition 2.3.1). Then,*

$$\mathbb{Z}^{D,LL,<\infty} \longrightarrow \mathbb{Z}^{LL,<\infty} \quad \textit{in probability as } |D| \to 0$$

where the limit is taken under the p-variation distance.

Part II

Signatures in finance

Chapter 3

Nonparametric pricing and hedging of exotic derivatives

This chapter is based on the paper [101], written in collaboration with Terry Lyons and Sina Nejad. In the spirit of Arrow-Debreu, we introduce a family of financial derivatives that act as primitive securities in that exotic derivatives can be approximated by their linear combinations. We call these financial derivatives signature payoffs. We show that signature payoffs can be used to nonparametrically price and hedge exotic derivatives in the scenario where one has access to price data for other exotic payoffs. The methodology leads to a computationally tractable and accurate algorithm for pricing and hedging using market prices of a basket of exotic derivatives that has been tested on real and simulated market prices, obtaining good results.

3.1 Introduction

Arrow-Debreu securities [5, 42] are *idealised* and *basic* securities that pay one unit of numeraire for a particular market state at a specific future time, and pay nothing otherwise. These securities are *primitive* in the sense that the cash flow of any derivative can be approximately written in terms of Arrow-Debreu securities. In this chapter, we identify a family of primitive securities for path-dependent exotic derivatives. Analogously to Arrow-Debreu securities, cash flows of exotic derivatives

can be approximated by linear combinations of these primitive securities.

When one has only limited access to market prices for a class of financial products, one may be interested in studying if knowledge of these prices can be leveraged to price other financial products in an arbitrage-free market [2, 47, 83]. For example, if one is able to price zero-coupon bonds, one could deduce the price of other coupon-bearing bonds by writing the cash flows as combinations of zero-coupon bonds. Similarly, if prices of European call and put options are observable in the market, it was shown in [17] that this information is enough to price any European contingent claim, by writing such contingent claims in terms of put and call options.

In this chapter, we take this idea a step further by showing that knowledge of prices of enough *exotic* derivatives is sufficient to accurately derive prices and hedging strategies of other *exotic* derivatives in a nonparametric manner. First, in the spirit of Arrow-Debreu [5, 42], we approximate exotic derivatives in terms of simpler payoffs called *signature payoffs* (Definition 3.2.6). Then, we estimate a certain quantity from the market: the *expected signature*. This procedure is empirically demonstrated in Section 3.5 and later in Chapter 4.

Signature payoffs are a family of path-dependent derivatives defined in terms of certain iterated integrals and, because of this, signature derivatives contain a lot of information about all possible dynamic trading strategies.

When one buys or sells a financial derivative on one or several assets, one is immediately exposed to certain risks. If this risk is unwanted, one may be interested in offsetting it by trading the underlying assets. This problem, known as *hedging*, is a classical problem in mathematical finance ([14, 109, 53]).

In idealised markets that are complete and frictionless, by definition it is possible to perfectly hedge such financial derivatives or payoffs. Therefore, in these cases the risk can theoretically be completely eliminated by following the hedging strategy.

However, transaction costs and other market frictions make real markets incomplete and hence it is not possible, in general, to perfectly hedge any given payoff. On top of that, market frictions such as transaction costs or liquidity constraints reduce the trader's ability to hedge. In these cases, one can try to find a hedging strategy that is optimal in the sense that it minimises a certain cost function. This cost function would be chosen by the trader, depending on her risk preferences.

An example of such optimisation problems is the mean-variance problem, where one wants to minimise the L^2 norm of the profits and losses (P&L) of the trading strategy ([126, 46, 125, 44]). This risk measure penalises any differences between the payoff and the corresponding hedging strategy. In particular, this risk preference penalises profits as well as losses. If one doesn't wish to penalise profits, the exponential utility function $x \mapsto \exp(-\lambda x)$ can be used instead, where $\lambda > 0$ is the risk-tolerance parameter ([43, 67]).

An open problem with obvious practical applications is how one could find a minimiser (or minimising sequence) for these optimal hedging problems. The paper [80] addresses this question for the mean-variance problem for vanilla options on Lévy processes, where the authors give a semi-explicit solution. This was later extended in [64], where the authors provide an algorithm for mean-variance hedging of vanilla options. In [85], on the other hand, the authors use BSDEs to try to get the optimal hedge for the mean-variance problem in the frictionless framework. Nevertheless, in general the optimal hedge seems to be difficult to find in practice. In [83], on the other hand, the authors use neural networks to price and hedge European options. In [13] this was extended to exotic derivatives, where the authors try to solve the optimisation problem by approximating the optimal hedging strategy with deep learning. However, this approach will, in general, only find local minima and not the global minima. Moreover, the training process can be computationally expensive.

In this chapter, we address the optimal hedging problem of minimising the expectation of a polynomial on the P&L. In particular, when the polynomial is chosen to be x^2, the classical mean-variance optimal hedging problem is retrieved. Stating the problem for general polynomials allows us to address the optimal hedging problem for the exponential utility function as well.

We make use of signatures from rough path theory ([105]) to reduce the general optimal hedging problem to a finite-dimensional optimisation problem that is computationally solvable. This is done in Section 3.3 by solving a linearised version of the problem first (see Theorem 3.3.3) and showing then that it suffices to solve this linearised problem to solve the original one (see Theorem 3.3.9). This method is then described in Algorithm 1. We assume that the price process X and volatility $\langle X \rangle$ of the underlying asset are such that $\mathbb{X}^{LL} := ((X, X), \langle X \rangle)$ is a 2-dimensional geometric rough path [105] which we call the lead-lag price path. It is shown in [52] that almost all sample paths X of continuous semimartingales, together with the quadratic variation $\langle X \rangle$, have this geometric rough path property. However, we do not impose any model on the dynamics of the asset, so that our approach is in this sense model-free.

The signature of a path is a transformation of path space that, in certain ways, behaves similarly to the Taylor expansion. Real-valued continuous functions on $\mathbb{R}^d$ are well approximated by linear functions on some polynomial basis. Similarly, real-valued continuous functions on some path space are well-approximated by linear functions on signatures (Proposition 3.3.6). This was already leveraged in the context of finance in [116], where payoffs were priced by writing them as linear functions on signatures. Signatures are concise and informative feature sets for paths, and as a consequence they have been used in machine learning in contexts other than finance, such as in mental health, handwriting recognition and gesture recognition ([65, 135, 96, 92, 141, 117]).

In Section 3.4 we solve several extensions of the original problem. For instance, in Sections 3.4.3 and 3.4.4 we solve the problem with market frictions – namely transaction costs and liquidity constraints. In Section 3.4.2, on the other hand, we study the semi-static hedging problem where the trader has access to a basket of derivatives for static hedging. Finally, in Section 3.4.5 we show how the optimal hedge can be found when the agent starts trading at a positive time after inception of the derivative.

Our approach is model-free in the sense that no model is assumed for the market dynamics. The only information needed from the market dynamics is its expected signature. When a particular model is imposed on the market dynamics, this expected signature can be computed using Monte Carlo methods or in certain cases by solving a PDE (see [112]). This approach is studied in Section 3.5, where we carry out some numerical experiments for a variety of payoffs and market models, obtaining good results. In Chapter 4 we will show that the expected signature can also be estimated in a model-free manner if market prices of sufficient exotic options are available.

3.2 Framework

3.2.1 The market

For simplicity, we will consider the case where there is only a single underlying risky asset. However, we would like to emphasise that all the results in this chapter can be readily extended to the multi-asset case. In the sequel, we will model the (discounted) price path of the underlying asset by a continuous curve in $\mathbb{R}$, $X : [0, T] \to \mathbb{R}$. We will denote the *augmentation* of X (as in [116]) by $\widehat{X}_t := (t, X_t) \in \mathbb{R}^2$. Without loss of generality, we will assume that the initial price of the asset is given by $X_0 = 1$. In our framework, the market will be given by the price path $\widehat{X} : [0, T] \to \mathbb{R}^2$, together with a volatility process $\langle \widehat{X} \rangle : [0, T] \to \mathbb{R}^{2\times 2}$. Almost all paths of a semimartingale

$X : [0,T] \to \mathbb{R}$ are included in this framework, in which case the volatility process $\langle \widehat{X} \rangle_t$ is just the quadratic variation of $\widehat{X}$. Tick-data is also included in this framework, as it induces a lead-lag path (see Section 2.3.1 and Chapter 4). However, our approach is model-free in the sense that we do not impose any model on the price path, nor do we assume it is a realisation of a semimartingale. Nevertheless, for simplicity the reader may think of X as a semimartingale. We will now introduce the precise definition of our market price paths.

Definition 3.2.1 (Market price paths)**.** *Define the space of market price paths,*

$$\widehat{\Omega}_T := \overline{\{\widehat{\mathbb{X}}^{<\infty} : X : [0,T] \to \mathbb{R} \text{ is smooth and } X_0 = 1\}}^{d_{p-var}} \subset T((\mathbb{R}^2)),$$

where $\widehat{X}_t := (t, X_t)$ denotes the augmentation of X, $\widehat{\mathbb{X}}^{<\infty}$ is the signature of $\widehat{X}$ and the closure is taken under the p-variation distance (Definition 2.2.1) for some $p \in [2,3)$. The space of lead-lag market price paths is defined by

$$\widehat{\Omega}_T^{LL} := \{\widehat{\mathbb{X}}^{LL,<\infty} : (\widehat{\mathbb{X}}^{\leq 2}, \langle \widehat{X} \rangle) \text{ is a lead-lag path and } X_0 = 1\} \subset T((\mathbb{R}^4)),$$

where $\widehat{\mathbb{X}}^{LL,<\infty}$ denotes the signature of the lead-lag path associated to $(\widehat{\mathbb{X}}^{\leq 2}, \langle \widehat{X} \rangle)$, as defined in Definition 2.3.1. Given $\widehat{\mathbb{X}}^{LL,<\infty} \in \widehat{\Omega}_T^{LL}$, we denote by $\widehat{\mathbb{X}}^{<\infty} \in \widehat{\Omega}_T$ the projection of $\widehat{\mathbb{X}}^{LL,<\infty}$ to $\widehat{\Omega}_T$.

So far, we have not imposed any probability measure on the market. We have only introduced the space of paths that will form the market, $\widehat{\Omega}_T^{LL}$. In Section 3.3 we will evaluate certain trading strategies by their performance in the market, for which we will use a probability measure. For this purpose, we will now define a probability space on the market paths. Most of the results in this chapter, however, are not dependent on the probability measure.

Definition 3.2.2. *Consider the Borel σ-algebra $\mathcal{B}(\widehat{\Omega}_T^{LL})$, and define by $\mathbb{F} = \{\mathcal{F}_t\}_{t\in[0,T]}$ the filtration generated by the price path X. Let $\mathbb{P}$ be a probability measure on*

$(\widehat{\Omega}_T^{LL}, \mathcal{B}(\widehat{\Omega}_T^{LL}))$ *such that* $\mathbb{E}[\widehat{\mathbb{X}}^{LL,\leq N}]$ *is finite for all* $N \in \mathbb{N}$. *We will then consider the completed filtered probability space* $(\widehat{\Omega}_T^{LL}, \mathcal{B}(\widehat{\Omega}_T^{LL}), \mathbb{F}, \mathbb{P})$.

We will not assume any particular model on the price – our approach is, in that sense, *model-free*. One could be interested, however, in imposing a particular model on the market, such as a certain semimartingale $X : [0, T] \to \mathbb{R}$. As discussed in Example 2.3.2, we can associate the semimartingale with the lead-lag path $(\widehat{\mathbb{X}}^{\leq 2}, \langle \widehat{X} \rangle)$ where $\langle X \rangle$ is the quadratic variation of X and $\widehat{\mathbb{X}}^{\leq 2}$ is the level-2 signature of $\widehat{X}$ (introduced in Example 2.2.5). Then, we would consider the probability space $(\widehat{\Omega}_T^{LL}, \mathcal{B}(\widehat{\Omega}_T^{LL}), \mathbb{F}, \mathbb{P})$ under which the coordinate process $X : [0, T] \to \mathbb{R}$ is a semimartingale with volatility given by the quadratic variation of $\widehat{X}$, i.e. $\langle \widehat{X} \rangle$.

Remark 3.2.3. The assumption that $\mathbb{E}[\widehat{\mathbb{X}}_{0,T}^{LL,\leq N}]$ exists for all $N \in \mathbb{N}$ is very mild, and it is an infinite-dimensional version of the "moments of all order exist" statement for finite-dimensional random variables.

3.2.2 Payoff functions

Financial derivatives are given in terms of a payoff function that depends on the underlying asset(s). We will now make a precise definition of a *payoff function*.

Definition 3.2.4 (Payoff function)**.** *A payoff is defined as a Borel-measurable function* $\widehat{\Omega}_T^{LL} \to \mathbb{R}$. *A payoff* $F : \widehat{\Omega}_T^{LL} \to \mathbb{R}$ *is said to be an* L^q*-payoff for* $q \geq 1$ *if* $\mathbb{E}[|F|^q] < \infty$.

The definition above essentially defines a payoff function as any $\mathcal{F}_T$-measurable random variable. The financial interpretation is that, given a realisation of the price path, the holder of the derivative with payoff $F : \widehat{\Omega}_T^{LL} \to \mathbb{R}$ is paid $F(\widehat{\mathbb{X}}^{LL,<\infty})$ at time T.

Example 3.2.5. Examples of payoff functions include European options, Asian options, lookback options, barrier options, futures, variance swaps, cliquet options, etc.

An important class of payoff functions, that will be used extensively in this chapter, are *linear signature payoff functions* ([116]):

Definition 3.2.6 (Linear signature payoff)**.** *We say that a payoff* $F : \widehat{\Omega}_T^{LL} \to \mathbb{R}$ *is a linear signature payoff function if there exists a linear functional* $f \in T((\mathbb{R}^4)^*)$ *such that*

$$F(\widehat{\mathbb{X}}^{LL,<\infty}) = \langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle.$$

These signature payoffs will play a similar role to Arrow-Debreu primitive securities. As we will see, path-dependent exotic payoffs can be well-approximated by these linear signature payoffs. Notice that because the signature is defined as certain iterated integrals against the path, linear signature payoffs effectively contain in particular the P&L of all dynamic hedging strategies. Therefore, in a way, it is unsurprising that the class of linear signature payoffs is *big* and that they form a family of primitive securities.

Example 3.2.7. We will now give a few examples of payoffs that can be written exactly as linear signature payoffs. Recall the word notation introduced in Section 1.3.

1. Let $K \in \mathbb{R}$, and set $f = (1-K)\varnothing + \mathbf{2}$. Then, $\langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle = 1 - K + X_T - X_0 = X_T - K$. In other words, the signature payoff is a forward with delivery price K.

2. Let $K \in \mathbb{R}$. Set $f = (1-K)\varnothing + \frac{1}{T}\mathbf{21}$. Then, $\langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle = 1 - K + \frac{1}{T}\int_0^T (X_s - X_0)ds = \frac{1}{T}\int_0^T X_s ds - K$. Therefore, Asian forwards are also signature payoffs.

3.2.3 Trading strategies

Intuitively, a trading strategy specifies the position that must be held by the trader at each time, given the observation of the price path up to that time. Moreover, this must be done in a non-anticipative way – in other words, traders are allowed to trade based on the past, but not the future. This idea is captured in the definition of trading strategies below.

Definition 3.2.8. *Define* $\Lambda_T := \bigcup_{t\in[0,T]} \widehat{\Omega}_t$, *which is a metric space for a certain distance. The space of trading strategies is defined by* $\mathcal{T}(\Lambda_T) := C(\Lambda_T;\mathbb{R})$. *We also denote by* $\mathcal{T}^q(\Lambda_T)$ *the space of trading strategies with the following integrability condition:*

$$\mathcal{T}^q(\Lambda_T) := \left\{\theta \in \mathcal{T}(\Lambda_T) : \mathbb{E}\left[\left|\int_0^T \theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,t]})dX_t\right|^q\right] < \infty\right\}.$$

Remark 3.2.9. The space Λ_T is the space of signatures of all *stopped paths*. A similar space was discussed in [36, 4, 57, 7] and in [48, 11, 120] in the context of finance.

Intuitively, the space of trading strategies from Definition 3.2.8 essentially consists of all non-anticipative processes with respect to the filtration generated by X. Again, this emphasises the crucial condition in finance that one is only allowed to trade based on the past.

We will now define an important subspace of the space of trading strategies – namely, the space of *linear signature* trading strategies. These are trading strategies that are given by linear combinations of the signature of the price path.

Definition 3.2.10 (Linear signature trading strategies)**.** *The space of linear signature trading strategies is given by*

$$\mathcal{T}_{sig}(\Lambda_T) := \{\theta \in \mathcal{T}(\Lambda_T) \mid \exists \ell \in T((\mathbb{R}^2)^*) \text{ such that } \theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,t]}) = \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle \ \forall\, \widehat{\mathbb{X}}|^{<\infty}_{[0,t]} \in \Lambda_T\}.$$

It turns out that, in some sense, trading strategies can be approximated arbitrarily well by signature trading strategies. Therefore, if one is looking for an optimal trading strategy in $\mathcal{T}(\Lambda_T)$, one could look for an optimal trading strategy in $\mathcal{T}_{sig}(\Lambda_T)$ instead. This will be made more precise later on.

Given a trading strategy $\theta \in \mathcal{T}$, the profits and losses (P&L) associated to it is given by the rough path integral (see [105]) given by $\int_0^T \theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,t]})dX_t$. In the particular case of semimartingales, this integral agrees with the classical Itô integral – see [52]. For instance, we have the following lemma, according to which Itô integrals of semimartingales of linear signature trading strategies are linear functions on the signature of the lead-lag path.

Lemma 3.2.11. *Let X be a d-dimensional continuous semimartingale. Let $\ell \in T((\mathbb{R}^2)^*)$. Then, we have:*

$$\int_0^T \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,t} \rangle dX_t = \langle \ell \mathbf{4}, \widehat{\mathbb{X}}^{LL,<\infty}_{0,T} \rangle,$$

where the integral is in the sense of Itô, the notation $\ell\mathbf{4} \in T((\mathbb{R}^4)^*)$ means the concatenation of the word associated to ℓ with the letter $\mathbf{4}$ (introduced in Section 1.3) and $\widehat{\mathbb{X}}^{LL,<\infty}_{0,T}$ is the signature of the (4-dimensional) lead-lag process, as defined in Definition 3.2.1 and Definition 2.3.1.

Therefore, the profit of a trading strategy, defined with a rough path integral, agrees with the classical definition in terms of an Itô integral in the case of semimartingales (see [52]) and in the particular case of linear signature trading strategies, the previous lemma states that the profits and losses – defined as an Itô integral against the semimartingale – is a linear functional of the signature of the lead-lag path.

3.3 Optimal hedging

In this section we study the following optimal polynomial hedging problem.

Definition 3.3.1 (Optimal polynomial hedging problem)**.** *Let $P \in \mathbb{R}[x]$ be a polynomial of degree $q \in \mathbb{N}$. Let F be an L^q-payoff that pays at terminal time T an amount of $F(\widehat{\mathbb{X}}^{LL,<\infty})$ with $\widehat{\mathbb{X}}^{LL,<\infty} \in \widehat{\Omega}_T^{LL}$. Let $p_0 \in \mathbb{R}$ be the initial capital. The associated optimal polynomial hedging problem (PHP) is to find a minimising sequence for:*

$$\inf_{\theta \in \mathcal{T}^q(\Lambda_T)} \mathbb{E}\left[P\left(F(\widehat{\mathbb{X}}^{LL,<\infty}) - p_0 - \int_0^T \theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,t]})dX_t\right)\right]. \tag{PHP}$$

In the particular case where $P(x) := x^2$, (PHP) is of the form of the well-studied mean-variance optimal hedging problem ([126, 46, 125, 44]). Writing the optimal control in terms of a polynomial P will allow us to extend (PHP) to the exponential utility function as well.

Our objective will be to provide a numerical approach to finding a minimising sequence for the optimal hedging problem. We will tackle (PHP) by studying a linearised version of the problem. As we will see in Section 3.3.2, solving this sub-problem will be sufficient to solve the original hedging problem (PHP).

3.3.1 A signature linearisation of the problem

Problem (PHP) will be solved by solving the following optimisation sub-problem instead:

Definition 3.3.2 (Optimal linear signature hedging problem)**.** *Let $P \in \mathbb{R}[x]$ be a polynomial of degree $q \in \mathbb{N}$. Let $f \in T((\mathbb{R}^4)^*)$ and consider the associated linear f-signature payoff (Definition 3.2.6). Define the optimal linear signature hedging problem (LSHP) as finding a minimising sequence for*

$$\inf_{\ell \in T((\mathbb{R}^2)^*)} \mathbb{E}\left[P\left(\langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle - p_0 - \int_0^T \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle dX_t\right)\right]. \tag{LSHP}$$

As we will see in Section 3.3.2, being able to solve (LSHP) will be sufficient to solve (PHP). (LSHP), on the other hand, can be rewritten as a simpler optimisation problem that is numerically easier to solve:

Theorem 3.3.3. *Let $f \in T((\mathbb{R}^4)^*)$ and $p_0 \in \mathbb{R}$. Let $P \in \mathbb{R}[x]$ be a polynomial of one variable, and recall the notation $P^{\text{ш}}$ introduced in Definition 1.3.5. Then, the solution of the optimal linear signature hedging problem* (LSHP) *is given by the solution of the following polynomial optimisation problem:*

$$\inf_{\ell \in T((\mathbb{R}^2)^*)} \left\langle P^{\text{ш}}(f - p_0 \varnothing - \ell \mathbf{4}), \mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\right] \right\rangle. \tag{3.1}$$

The optimisation problem (3.1) has two components. First, there is a linear functional, which depends on the control $\ell \in T((\mathbb{R}^2)^*)$ over which we are optimising, but does not depend on the price path X. The second component is the expected signature of the lead-lag process $\mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\right]$, which clearly depends on the price path X, but does not depend on the control. Notice that if a risk-neutral measure is used instead of the real-world probability measure, knowing the expected signature is equivalent of knowing the prices of all signature payoffs.

So far we haven't made any assumptions on the price path X nor the volatility $\langle X \rangle$. Hence, the only information one needs about the process to optimally hedge an exotic derivative is the expected signature of the lead-lag path corresponding to the underlying asset. If one wants to assume a particular model for the price process (such as an Itô process or semimartingale), the expected signature can be computed either using Monte Carlo methods or, in certain cases, by solving a PDE. A more detailed discussion about this will be made in Section 3.3.3. In the next chapter we will show that if the trader has access to sufficient market prices of exotic options, this extra information is enough to estimate the expected signature (called *implied expected signature*) in a model-free manner.

A consequence of Theorem 3.3.3 is the following corollary, which gives sufficient conditions for a linear signature payoff to be *attainable* or *replicable*.

Corollary 3.3.4. *Let $f \in T((\mathbb{R}^4)^*)$. Assume that f is of the form*

$$f = p_0 \varnothing + \sum_{n=0}^{N} \sum_{\substack{\mathbf{w}=\mathbf{i_1}\dots\mathbf{i_n} \\ i_j \in \{1,2\}}} \lambda_{\mathbf{w}} \mathbf{i_1} \dots \mathbf{i_n} \mathbf{4}$$

where $N \in \mathbb{N}$ and $p_0, \lambda_{\mathbf{w}} \in \mathbb{R}$. Then, the linear signature payoff given by f is attainable and the optimal hedging strategy is given by the linear signature trading strategy

$$t \mapsto \sum_{n=0}^{N} \sum_{\substack{\mathbf{w}=\mathbf{i_1}\dots\mathbf{i_n} \\ i_j \in \{1,2\}}} \lambda_{\mathbf{w}} \left\langle \mathbf{i_1} \dots \mathbf{i_n}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle.$$

Example 3.3.5. The Asian-style payoff given by $F(\widehat{\mathbb{X}}^{LL,<\infty}) := X_T - \frac{1}{T}\int_0^T X_t dt$ is replicable and the replicating strategy is given by $\theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,t]}) = t/T$. Indeed, we have:

$$\begin{aligned} F(\widehat{\mathbb{X}}^{LL,<\infty}) &= \langle \mathbf{4}, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle - \frac{1}{T} \langle \mathbf{41}, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle \\ &= \frac{1}{T} \langle \mathbf{1}, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle \langle \mathbf{4}, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle - \frac{1}{T} \langle \mathbf{41}, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle \\ &= \frac{1}{T} \langle \mathbf{1} \sqcup\!\sqcup \mathbf{4} - \mathbf{41}, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle = \frac{1}{T} \langle \mathbf{14}, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle \end{aligned}$$

and the previous corollary applies.

3.3.2 The general optimal hedging problem

The optimisation problem (3.1) offers an implementable way of computing the optimal hedge in (LSHP) – for example, in the case of mean-variance hedging where $P(x) := x^2$, the optimisation problem (3.1) is reduced to finding the global minimum of a high-dimensional quadratic polynomial, which in turn can be easily found by solving a certain system of linear equations (see Section 3.3.3 for a discussion on numerically solving (LSHP) in practice).

However, the ultimate goal is to solve the optimal hedging problem shown in (PHP). In what follows, we justify why we may replace a general payoff function

$F : \widehat{\Omega}_T^{LL} \to \mathbb{R}$ with a signature payoff, and why we may restrict the class of trading strategies from $\mathcal{T}^q(\Lambda_T)$ to $\mathcal{T}_{sig}(\Lambda_T)$. This section culminates in Theorem 3.3.9, which allows us to consider the tractable optimal hedging problem (LSHP) (and hence (3.1)) instead of the original nonlinear, hard-to-solve problem (PHP).

See Appendix B for the detailed justification with proofs of why solving (LSHP) is enough to solve (PHP).

From payoffs to signature payoffs

Arrow-Debreu securities are securities that pay 1 if a certain state of the market occurs, and nothing otherwise. Hence, exotic derivatives can be decomposed as linear combinations of such securities – Arrow-Debreu securities are, in other words, *primitive securities* from which all other securities are built.

In a similar fashion, the iterated integrals that define signatures (Definition 2.1.1) are primitive securities in the sense that other exotic, path-dependent payoffs are well-approximated by linear combinations of such iterated integrals. They are, effectively, basic securities from which other securities are built. Moreover, given that signature payoffs are defined as linear combinations of certain iterated integrals, they include a lot of information about the P&L of all possible dynamic trading strategies.

This is made precise in the following proposition and corollary.

Proposition 3.3.6. *Let $F : \widehat{\Omega}_T^{LL} \to \mathbb{R}$ be a continuous payoff and let $\mathcal{K} \subset \widehat{\Omega}_T^{LL}$ be a compact set. Given any $\varepsilon > 0$, there exists $f \in T((\mathbb{R}^4)^*)$ such that*

$$|F(\widehat{\mathbb{X}}^{LL,<\infty}) - \langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}^{LL,<\infty} \in \mathcal{K}.$$

Corollary 3.3.7. *Let $F : \widehat{\Omega}_T^{LL} \to \mathbb{R}$ be a continuous payoff. Given any $\varepsilon > 0$, there exists a compact set $\mathcal{K}_\varepsilon \subset \widehat{\Omega}_T$ (which does not depend on F) and $f \in T((\mathbb{R}^4)^*)$ such that:*

1. $\mathbb{P}[\mathcal{K}_\varepsilon] > 1 - \varepsilon$,

2. $|F(\widehat{\mathbb{X}}^{LL,<\infty}) - \langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}^{LL,<\infty} \in \mathcal{K}_\varepsilon$.

In other words, there exists a large compact set – large in the sense that with very high probability, all price paths one observes lie on that compact set – such that on the compact set, all continuous payoffs look like signature payoffs.

The author wants to emphasise that the linear functional f from Corollary 3.3.7 does not depend on any model for the underlying assets. Indeed, it is a pathwise and a model-free density result that does not require any probability structure. The only role of the probability measure $\mathbb{P}$ in Corollary 3.3.7 is providing a notion of *big* compact sets – i.e. point 1. in Corollary 3.3.7.

From trading strategies to signature trading strategies

In Section 3.2.3, we defined the space of trading strategies $\mathcal{T}$ – which intuitively consists of all adapted processes – as well as the subspace of signature trading strategies $\mathcal{T}_{sig} \subset \mathcal{T}$. Similarly to signature payoffs, the space of signature trading strategies is big in the sense that arbitrary trading strategies can be well-approximated by them:

Proposition 3.3.8. *Let $\mathcal{K} \subset \Lambda_T$ be a compact set. Then, given any trading strategy $\theta \in \mathcal{T}$, there exists $\ell \in T((\mathbb{R}^2)^*)$ such that*

$$|\theta(\widehat{\mathbb{X}}|_{[0,t]}^{<\infty}) - \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}|_{[0,t]}^{<\infty} \in \mathcal{K}.$$

Putting everything together: from optimal hedging to optimal signature hedging

A consequence of Corollary 3.3.7 and Proposition 3.3.8 is the following theorem, which justifies why we can consider the optimal linear signature problem (LSHP) instead of the original optimal hedging problem (LSHP).

Theorem 3.3.9. *Let*

$$a := \inf_{\theta \in \mathcal{T}^q(\Lambda_T)} \mathbb{E}\left[P\left(F(\widehat{\mathbb{X}}^{LL,<\infty}) - p_0 - \int_0^T \theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,t]})dX_t\right)\right]$$

be the infimum of the optimal polynomial hedging problem (PHP). *Given any $\varepsilon > 0$, there exists a compact set $\mathcal{K}_\varepsilon \subset \widehat{\Omega}_T$, a linear signature payoff given by $f \in T((\mathbb{R}^4)^*)$ and a linear signature trading strategy given by $\ell \in T((\mathbb{R}^2)^*)$ such that:*

1. $\mathbb{P}[\mathcal{K}_\varepsilon] > 1 - \varepsilon$,
2. $|F(\widehat{\mathbb{X}}^{LL,<\infty}) - \langle f, \widehat{\mathbb{X}}^{LL,<\infty}_{0,T}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}^{LL,<\infty} \in \mathcal{K}_\varepsilon$,
3. $|\theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,t]}) - \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}^{<\infty} \in \mathcal{K}_\varepsilon$ *and* $t \in [0,T]$,
4. $|a_\varepsilon - a| \leq \varepsilon$, *where*

$$a_\varepsilon := \mathbb{E}\left[P\left(\langle f, \widehat{\mathbb{X}}^{LL,<\infty}_{0,T}\rangle - p_0 - \int_0^T \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle dX_t\right) \;;\; \mathcal{K}_\varepsilon\right].$$

That is, being able to solve the optimal linear signature hedging problem (LSHP) provides a numerically feasible way of finding a minimising sequence of the optimal polynomial hedging problem (PHP).

3.3.3 Solving the optimal linear signature hedging problem

Theorem 3.3.3, together with Theorem 3.3.9, offers an implementable approach to numerically approximating the optimal hedging strategy for a polynomial P of degree q, an L^q-payoff F and initial cash p_0.

The only information that is needed about the process is its expected signature. The expected signature plays a role similar to the moments of a real-valued random variable, but on path space – i.e. under certain growth assumptions, the expected signature of a process determines the law of the process [33]. Therefore, the fact that

Algorithm 1: Estimating the optimal hedge.

Parameters: $T > 0$: terminal time.
$P \in \mathbb{R}[x]$: polynomial of degree $q \in \mathbb{N}$.
F: an L^q payoff.
$p_0 \in \mathbb{R}$: initial capital.
$N \geq 2$: order of the signature.
$\widehat{\mathbb{X}}^{LL,<\infty}$: market price path.

Output: An estimation $\ell \in T^{(\lfloor N/q \rfloor)}((\mathbb{R}^2)^*)$ of the optimal hedge.

1 Take a finite training dataset $\mathcal{D} \subset \widehat{\Omega}_T$.
2 Transform $\mathcal{D}$ into a dataset of truncated signatures of order N, i.e. $\mathcal{D}_N := \{\widehat{\mathbb{Y}}_{0,T}^{LL,\leq N} : \widehat{\mathbb{Y}}^{LL,<\infty} \in \mathcal{D}\}$.
3 Compute the payoffs $F(\mathcal{D}) \subset \mathbb{R}$.
4 Apply linear regression to $\mathcal{D}_N$ against $F(\mathcal{D})$ to find $f \in T^{(N)}((\mathbb{R}^4)^*)$ such that $\langle f, \widehat{\mathbb{Y}}_{0,T}^{LL,\leq N} \rangle \approx F(\widehat{\mathbb{Y}}^{LL,<\infty})$ for each $\widehat{\mathbb{Y}}^{LL,<\infty} \in \mathcal{D}$.
5 Estimate the expected signature $\mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{LL,\leq N}\right]$.
6 Find a minimiser ℓ of the optimisation problem

$$\begin{aligned} &\text{minimise} \quad \left\langle P^{ш}\left(f - p_0 \varnothing - \ell \mathbf{4}\right), \mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{\mathrm{LL},\leq N}\right] \right\rangle \\ &\text{over} \quad \ell \in T^{(\lfloor N/q \rfloor)}((\mathbb{R}^2)^*). \end{aligned}$$

7 **return** ℓ.

the optimal hedge depends on the expected signature only essentially means that it depends on the whole law of the dynamics of the price path.

For obvious computational reasons, one cannot work with the whole expected signature – one has to begin by fixing a signature order $N \in \mathbb{N}$ and considering the corresponding truncated signature of the lead-lag process $\mathbb{E}[\widehat{\mathbb{X}}_{0,T}^{LL,\leq N}]$. This truncated expected signature can be computed using Monte Carlo methods or even by solving a certain PDE (see [112]) in the case when one imposes a model on the market dynamics. We will see in Chapter 4 that the trader can also estimate the expected signature in a model-free way if the trader has access to market prices of sufficient exotic options.

Remark 3.3.10. The shuffle product of a word of length n and a word of length

m is a sum of words of lengths $m+n$. Therefore, when considering the truncated signature $\mathbb{E}[\widehat{\mathbb{X}}_{0,T}^{LL,\leq N}]$ and a polynomial P of degree q, (3.1) has to be minimised over $\ell \in T^{(\lfloor N/q \rfloor)}((\mathbb{R}^2)^*)$ rather than $\ell \in T((\mathbb{R}^2)^*)$.

Once (3.1) is restricted to a signature order $N \in \mathbb{N}$, the optimisation problem (3.1) is reduced to finding the minimum of a high-dimensional polynomial of degree q. In particular, in the mean-variance optimal hedging problem where the polynomial is given by $P(x) := x^2$, solving (3.1) consists of solving a system of linear equations.

Algorithm 1 describes the proposed algorithm. We make some practical remarks. Signatures can be computed using the publicly available software `esig`[1]. An alternative package is `iisignature`[2] or `signatory`[3]. Finally, the recursive definition of the shuffle product in Definition 1.3.3 allows for an easy and efficient implementation.

Once we have found a minimising linear functional $\ell \in T^{(\lfloor N/q \rfloor)}((\mathbb{R}^2)^*)$, the hedging strategy would be given by

$$t \mapsto \langle \ell, \widehat{\mathbb{X}}_{0,t}^{\leq N} \rangle \quad \forall t \in [0,T].$$

3.4 Extensions

In this section we will discuss a few extensions of our original framework. Needless to say, these extensions can be combined depending on the features one wishes the optimal hedging problem to have.

3.4.1 Exponential hedging

Often, one is interested in only hedging unfavourable differences between the hedged derivative and the hedging strategy. In other words, one wants to penalise losses and reward profits. This can be accomplished by considering the exponential hedging

[1] `https://pypi.org/project/esig/`
[2] [119], `https://github.com/bottler/iisignature`
[3] [88], `https://github.com/patrick-kidger/signatory`

problem, where instead of minimising the expectation of a polynomial on the P&L, one replaces the polynomial by an exponential function $x \mapsto \exp(-\lambda x)$ for some risk parameter $\lambda > 0$.

We will begin by introducing the space of trading strategies that we will consider admissible.

Definition 3.4.1 (Admissible trading strategy for exponential hedging)**.** *Define the space of admissible trading strategies*

$$\mathcal{T}^{\infty}(\Lambda_T) := \left\{ \theta \in \mathcal{T}(\Lambda_T) : \int_0^T \theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,u]}) dX_u \text{ is bounded a.s.} \right\}.$$

We may now define the optimal exponential hedging problem as follows:

Definition 3.4.2 (Optimal exponential hedging problem)**.** *Let $\lambda > 0$. Let F be a payoff that is bounded a.s. Let $p_0 \in \mathbb{R}$. The associated optimal exponential problem is:*

$$\inf_{\theta \in \mathcal{T}^{\infty}(\Lambda_T)} \mathbb{E}\left[\exp\left(-\lambda \left(p_0 + \int_0^T \theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,t]}) dX_t - F(\widehat{\mathbb{X}}^{LL,<\infty}) \right) \right) \right] \tag{3.2}$$

The parameter $\lambda > 0$ measures the risk tolerance of the trader: the greater it is, the less tolerant the trader is with respect to losses.

The following proposition shows that we can tackle the optimal hedging problem (3.2) by reducing it to an optimal hedging problem of the form (LSHP), which was solved in Section 3.3.

Proposition 3.4.3. *Let*

$$a := \inf_{\theta \in \mathcal{T}^{\infty}(\Lambda_T)} \mathbb{E}\left[\exp\left(-\lambda \left(p_0 + \int_0^T \theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,t]}) dX_t - F(\widehat{\mathbb{X}}^{LL,<\infty}) \right) \right) \right]$$

be the infimum of the optimal exponential hedging problem. Given any $\varepsilon > 0$, there exists a polynomial $P_\varepsilon \in \mathbb{R}[x]$, a compact set $\mathcal{K}_\varepsilon \subset \widehat{\Omega}_T$, a linear signature payoff given by $f \in T((\mathbb{R}^4)^*)$ and a linear signature trading strategy given by $\ell \in T((\mathbb{R}^2)^*)$ such that:

1. $P_\varepsilon \xrightarrow{\varepsilon\to 0} \exp(-\lambda\,\cdot)$ *uniformly on compact intervals,*

2. $\mathbb{P}[\mathcal{K}_\varepsilon] > 1-\varepsilon$,

3. $|F(\widehat{\mathbb{X}}^{LL,<\infty}) - \langle f, \widehat{\mathbb{X}}^{LL,<\infty}_{0,T}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}^{LL,<\infty} \in \mathcal{K}_\varepsilon$,

4. $|\theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,t]}) - \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}^{<\infty} \in \mathcal{K}_\varepsilon$ *and* $t \in [0,T]$,

5. $|a_\varepsilon - a| \le \varepsilon$, *where*

$$a_\varepsilon := \mathbb{E}\left[P_\varepsilon\left(p_0 + \int_0^T \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle dX_t - \langle f, \widehat{\mathbb{X}}^{LL,<\infty}_{0,T}\rangle\right)\ ;\ \mathcal{K}_\varepsilon\right].$$

3.4.2 Semi-static hedging

In certain situations, one wants to hedge a payoff $F : \widehat{\Omega}^{LL}_T \to \mathbb{R}$, for which one has access to a (finite) basket of derivatives $\mathcal{B} = \{G_i\}_{i=0}^k$ that are allowed for static hedging, as well as the underlying asset X that can be used for dynamic hedging. In other words, at inception $t = 0$ the trader has to form a portfolio on the basket $\mathcal{B}$ and then the trader can dynamically trade on the underlying X. For example, the payoff F we wish to hedge could be an exotic option, and the basket $\mathcal{B}$ could be a basket of simpler vanilla options. We will assume that the payoff F, as well as each payoff $G_i \in \mathcal{B}$, are L^q payoffs. Moreover, following Section 3.3.1 and Section 3.3.2, we will assume that F is a f-signature payoff and G_i is a g_i-signature payoff, with $f, g_i \in T((\mathbb{R}^4)^*)$.

If we allow semi-static hedging, the optimal hedging problem is then defined as:

$$\inf_{\substack{\ell\in T((\mathbb{R}^2)^*)\\(\beta_i)_{i=1}^k\in\Gamma}} \mathbb{E}\left[P\left(\langle f, \widehat{\mathbb{X}}^{LL,<\infty}_{0,T}\rangle - p_0 - \sum_{i=1}^k \beta_i \langle g_i, \widehat{\mathbb{X}}^{LL,<\infty}_{0,T}\rangle - \int_0^T \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle dX_t\right)\right]$$

where $\Gamma \subset \mathbb{R}^k$ determines the region of admissible strategies on $\mathcal{B}$. For example, if no constraints are imposed, one could choose $\Gamma = \mathbb{R}^k$. If no short-selling is allowed, on the other hand, one would choose $\Gamma = \mathbb{R}^k_+$. Other choices are also allowed, which

can include liquidity constraints or other more complex, inter-connected constraints on $\mathcal{B}$.

Then, the semi-static optimal hedging problem is then reduced to the following.

Corollary 3.4.4. *Given a basket of signature payoffs $\mathcal{B} = \{g_i\}_{i=1}^k$ and a region $\Gamma \subset \mathbb{R}^k$, the solution of the semi-static optimal hedging problem is given by the solution of*

$$\inf_{\substack{\ell \in T((\mathbb{R}^2)^*) \\ (\beta_i)_{i=1}^k \in \Gamma}} \left\langle P^{\text{ш}} \left(f - \sum_{i=1}^k \beta_i g_i - p_0 \varnothing - \ell \mathbf{4} \right), \mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \right] \right\rangle.$$

3.4.3 Adding transaction costs

It will turn out, unsurprisingly, that linear signature trading strategies with wild oscillations (i.e. infinite total variation) will incur infinite transaction costs. This could be avoided by considering linear signature trading strategies that have a well-defined *speed of trading*:

$$\begin{aligned} \mathcal{T}_{speed}(\Lambda_T) :&= \left\{ \Lambda_T \ni \widehat{\mathbb{X}}|^{<\infty}_{[0,t]} \mapsto \int_0^t \langle v, \widehat{\mathbb{X}}^{<\infty}_{0,u} \rangle du \mid v \in T((\mathbb{R}^2)^*) \right\} \\ &= \left\{ \Lambda_T \ni \widehat{\mathbb{X}}|^{<\infty}_{[0,t]} \mapsto \langle v\mathbf{1}, \widehat{\mathbb{X}}^{<\infty}_{0,t} \rangle \mid v \in T((\mathbb{R}^2)^*) \right\}. \end{aligned}$$

The function $\langle v, \widehat{\mathbb{X}}^{<\infty}_{0,t} \rangle$ indicates the trading speed – i.e. the amount of the underlying asset that will be bought or sold at each time infinitesimal time. For such a choice of trading speed, the trader's position on the asset at time t will be $\int_0^t \langle v, \widehat{\mathbb{X}}^{<\infty}_{0,u} \rangle du$. Therefore, such trading strategies will be differentiable and hence they will not incur infinite transaction costs. This approach to defining trading speeds has been extensively used in the literature, i.e. [22, 24, 93, 23].

We will now introduce the following *fixed* and *proportional* quadratic costs for trading strategies in $\mathcal{T}_{speed}$. As noted by [31], transaction costs per unit traded may be linear in the traded amount, suggesting that quadratic transaction costs might be appropriate. Intuitively, these transaction costs penalise large transactions heavily,

while comparatively smaller transactions have little cost. Quadratic transaction costs can also be justified from a market microstructure point of view through a *temporary market impact*, see Chapter 5.

Definition 3.4.5 (Fixed quadratic transaction costs)**.** *Consider a speed of trading $v \in T((\mathbb{R}^2)^*)$. We define the fixed quadratic costs with intensity $\alpha \geq 0$ incurred by v along $\widehat{\mathbb{X}}^{LL,<\infty} \in \widehat{\Omega}_T$ as*

$$C_\alpha^{\text{fixed}}(v, \widehat{\mathbb{X}}^{LL,<\infty}) := \alpha \int_0^T |\langle v, \widehat{\mathbb{X}}_{0,u}^{<\infty} \rangle|^2 du.$$

Definition 3.4.6 (Proportional quadratic transaction costs)**.** *Consider a speed of trading $v \in T((\mathbb{R}^2)^*)$. The proportional quadratic costs with intensity $\alpha \geq 0$ incurred by v along $\widehat{\mathbb{X}}^{LL,<\infty} \in \widehat{\Omega}_T$ is then defined as*

$$C_\alpha^{\text{prop}}(v, \widehat{\mathbb{X}}^{LL,<\infty}) := \alpha \int_0^T |\langle v, \widehat{\mathbb{X}}_{0,u}^{<\infty} \rangle X_u|^2 du.$$

Then, one can naturally modify (LSHP) to include *fixed* quadratic costs,

$$\inf_{v \in T((\mathbb{R}^2)^*)} \mathbb{E}\left[P\left(\langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle - p_0 - \int_0^T \int_0^t \langle v, \widehat{\mathbb{X}}_{0,u}^{<\infty} \rangle du \, dX_t + C_\alpha^{\text{fixed}}(v, \widehat{\mathbb{X}}^{LL,<\infty}) \right) \right], \tag{3.3}$$

or *proportional* transaction costs,

$$\inf_{v \in T((\mathbb{R}^2)^*)} \mathbb{E}\left[P\left(\langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle - p_0 - \int_0^T \int_0^t \langle v, \widehat{\mathbb{X}}_{0,u}^{<\infty} \rangle du \, dX_t + C_\alpha^{\text{prop}}(v, \widehat{\mathbb{X}}^{LL,<\infty}) \right) \right]. \tag{3.4}$$

We then have the following corollaries of Theorem 3.3.3.

Corollary 3.4.7. *The solution of the optimal hedging problem under fixed quadratic trading costs* (3.3) *is given by the solution of the following optimisation problem:*

$$\inf_{v \in T((\mathbb{R}^2)^*)} \left\langle P^{\text{ш}}(f - p_0 \varnothing - v\mathbf{14} + \alpha v^{\text{ш}2}\mathbf{1}), \mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \right] \right\rangle. \tag{3.5}$$

Similarly, the solution of the optimal hedging under proportional transaction costs (3.4) *is given by*

$$\inf_{v \in T((\mathbb{R}^2)^*)} \left\langle P^{\text{ш}}(f - p_0 \varnothing - v\mathbf{14} + \alpha(v\text{ш}(\mathbf{2} + \varnothing))^{\text{ш}2}\mathbf{1}), \mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \right] \right\rangle.$$

3.4.4 Liquidity constraints

So far, we have implicitly assumed that the market for the underlying asset is perfectly liquid: we can, at any given time, take a long or short position of any size. However, this assumption may not be realistic for some assets, so that one has to impose certain liquidity constraints. We will do so by imposing a certain boundedness condition on the speed of trading that was introduced in Section 3.4.3. More specifically, we will consider all trading speeds $v \in T((\mathbb{R}^2)^*)$ such that $\|v\| \leq M$, for some *illiquidity constant* $M \geq 0$. This parameter could be estimated from historical data of the asset, for example.

Under liquidity constraints, the unconstrained optimal hedging problem (LSHP) is transformed to the following:

Definition 3.4.8 (Optimal hedging problem with liquidity constraints)**.** *Given an illiquidity constant $M \geq 0$, the following problem is defined as the optimal hedging strategy with liquidity constraint M:*

$$\inf_{\substack{v \in T((\mathbb{R}^2)^*) \\ \|v\| \leq M}} \mathbb{E}\left[P\left(\langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle - p_0 - \int_0^T \int_0^t \langle v, \widehat{\mathbb{X}}_{0,u}^{<\infty}\rangle du\, dX_t\right)\right]. \tag{3.6}$$

The solution of the optimal hedging problem with liquidity constraints is then a constrained optimisation problem:

Corollary 3.4.9. *Let $M \geq 0$ be an illiquidity constant. The solution of the optimal hedging problem with liquidity constraint M is given by*

$$\inf_{\substack{v \in T((\mathbb{R}^2)^*) \\ \|v\| \leq M}} \left\langle P^{\text{ш}}(f - p_0 \varnothing - v\mathbf{14}), \mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\right]\right\rangle.$$

3.4.5 Delayed hedging

Suppose a trader wishes to hedge a certain derivative given by the payoff $F : \widehat{\Omega}_T^{LL} \to \mathbb{R}$, whose lifespan is $[0, T]$. However, the trader is at time $t > 0$, so that she cannot follow

the hedging strategy provided by Theorem 3.3.3 as the trader will only be able to trade over $[t, T]$. What is the optimal strategy that the trader can carry?

Again, by Section 3.3.2 we will assume that F is a f-signature payoff for $f \in T((\mathbb{R}^4)^*)$. The objective is then to solve the following optimal hedging problem:

$$\inf_{\ell \in T((\mathbb{R}^2)^*)} \mathbb{E}\left[P\left(\langle f, \widehat{\mathbb{X}}_{0,T}^{<\infty} \rangle - p_t - \int_t^T \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle dX_t \right) \Big| \mathcal{F}_t \right] \tag{3.7}$$

where p_t is $\mathcal{F}_t$-measurable and represents the cash held at time t. Notice that

$$\begin{aligned} \int_t^T \langle \ell, \widehat{\mathbb{X}}_{0,u}^{<\infty} \rangle dX_u &= \int_0^T \langle \ell, \widehat{\mathbb{X}}_{0,u}^{<\infty} \rangle dX_u - \int_0^t \langle \ell, \widehat{\mathbb{X}}_{0,u}^{<\infty} \rangle dX_u \\ &= \langle \ell \mathbf{4}, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle - \langle \ell \mathbf{4}, \widehat{\mathbb{X}}_{0,t}^{LL,<\infty} \rangle. \end{aligned}$$

Moreover, notice that $\langle \ell \mathbf{4}, \widehat{\mathbb{X}}_{0,t}^{LL,<\infty} \rangle \in \mathbb{R}$ is known at time t. We then have:

Corollary 3.4.10. *The solution of the optimal hedging problem* (3.7) *where the trader starts hedging at time $t > 0$ is given by*

$$\inf_{\ell \in T((\mathbb{R}^2)^*)} \left\langle P^{\text{ш}} \left(f - \left(p_t + \langle \ell \mathbf{4}, \widehat{\mathbb{X}}_{0,t}^{\mathrm{LL},<\infty} \rangle \right) \varnothing - \ell \mathbf{4} \right), \widehat{\mathbb{X}}_{0,t}^{\mathrm{LL},<\infty} \otimes \mathbb{E}^{\mathbb{P}} \left[\widehat{\mathbb{X}}_{t,T}^{\mathrm{LL},<\infty} \Big| \mathcal{F}_t \right] \right\rangle. \tag{3.8}$$

In other words, all the trader has to do is compute the signature up to time t of the lead-lag process of the augmented price path, as well as the expected signature for the remaining interval $[t, T]$, and then solve the optimisation problem (3.8).

3.5 Experiments on synthetic data

Our methodology is intrinsically nonparametric, as the only information required from the market is its expected signature. In this section, we will impose a model on the market which is known to the trader. This situation is common in industry – for example, banks are often interested in pricing and hedging exotic payoffs when

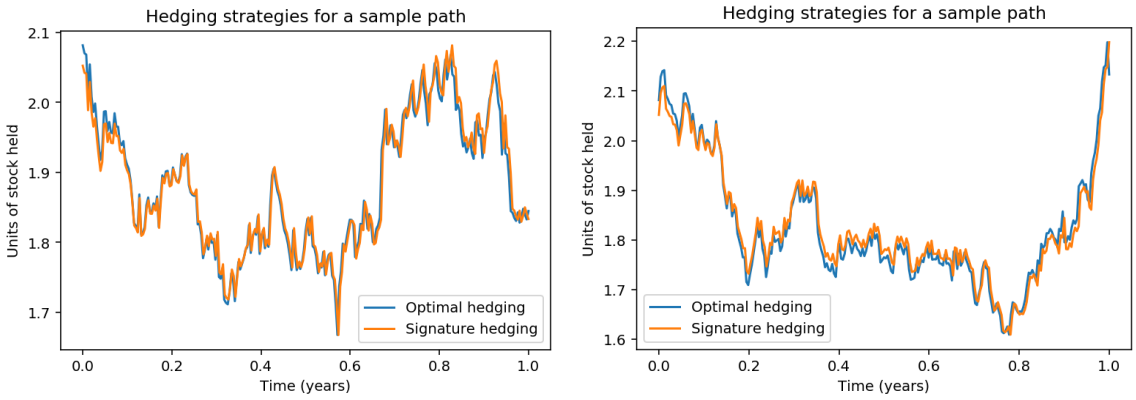

Figure 3.1: Units of stock held by the signature hedging strategy and the optimal hedging strategy on two realizations of the Black–Scholes model.

the market dynamics are given by one of the internal models of the bank. In the next chapter, we will show that if the market dynamics are unknown – which is what happens in real markets – the trader can still estimate the expected signature (and therefore price and hedge exotic options) in a model-free way provided the trader has access to market prices of other vanilla and exotic options.

In the remainder of this section, we implement the proposed approach in a wide range of examples to show the effectiveness of the methodology on different market models. We will begin by considering in Section 3.5.1 a toy example with a simple payoff in a complete market, in order to compare the signature hedging strategy with the (known) replicating strategy. Then, in Section 3.5.2 we will implement our methodology on path-dependent payoffs in an incomplete market. In Section 3.5.3 we will consider the exponential hedging problem and finally in Section 3.5.4 we will study the hedging problem under transaction costs.

3.5.1 Toy example

First, we considered the simple case where we assume that X follows a Black–Scholes model and we want to hedge the derivative with payoff $F(\widehat{\mathbb{X}}^{LL,<\infty}) = X_T^2$ at terminal time T. This payoff, under this model, is attainable and we should therefore be able

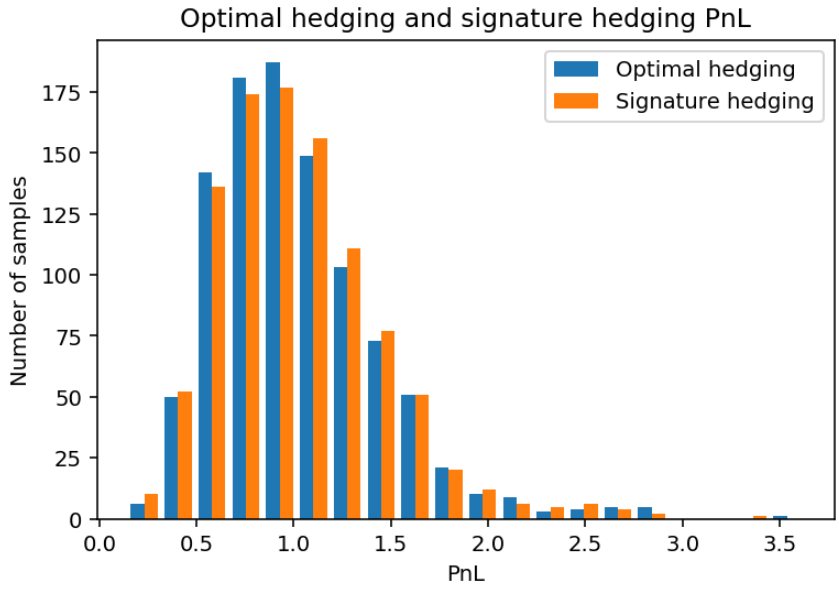

Figure 3.2: P&L of the replicating and signature hedging strategies for the toy example in Section 3.5.1.

to perfectly hedge it. Given that we can explicitly find what the replicating strategy should be – it is the delta hedge – this example will be useful to determine whether the optimal signature strategy matches the replicating strategy.

We implemented Algorithm 1 for the polynomial $P(x) := x^2$, so that we are considering the mean-variance hedging problem. The initial capital p_0 was taken to be the risk-neutral price for the payoff, and the signature order was set to 8. We fixed the maturity to 1 year ($T = 1$), and we assumed daily rebalancing.

Figure 3.1 shows the number of units of stock held by the optimal continuous-time hedging strategy (i.e. the replicating strategy) and the signature hedging strategy provided by Algorithm 1, for two realisations of the underlying price process. As we see, both strategies match very well. Note that rebalancing occurs daily in this experiment, not continuously, so the optimal continuous-time hedging strategy will not be optimal in general when trading is done discretely.

Consider the P&L of the replicating and signature hedging strategies in Figure 3.2, which should (almost) match the P&L of the option. We observe that both

strategies have very similar P&L distributions.

3.5.2 Path-dependent payoffs on the Heston model

We now consider path-dependent payoffs on the Heston model, which is incomplete. The payoffs we considered were Asian options, barrier call options, lookback options and variance swaps. As in the previous section, we considered the mean-variance hedging problem with maturity 1 year, and we set p_0 to be the risk-neutral price for each payoff. Again, the signature order we considered was 8.

Figure 3.3 shows the P&L of the hedged portfolio at maturity with daily rebalancing. Ideally, the payoffs would be perfectly hedged so that the P&L of the hedged portfolio would be identically zero. However, given that the Heston model is incomplete and we are considering daily rebalancing, this is not possible in general.

3.5.3 Exponential hedging

If we change the risk preferences of the trader in order to penalise losses but not profits, we may consider the exponential hedging problem rather than the mean-variance hedging problem. Following Section 3.4.1, we approximate $x \mapsto \exp(-\lambda x)$ by polynomials. We then solve the optimal linear signature hedging problem for the Asian option payoff. Notice that this payoff is not a.s. bounded and it therefore does not satisfy the hypotheses of Proposition 3.4.3. However, we can overcome this issue by assuming that the payoff was truncated on $[-M, M]$, for $M > 0$ large enough.

The performance of the signature hedging strategy is shown in Figure 3.4, where the risk parameter $\lambda = 0.25$ was considered. Notice that this risk parameter has shifted the P&L profile from Figure 3.3a, reflecting the change in the trader's risk preferences.

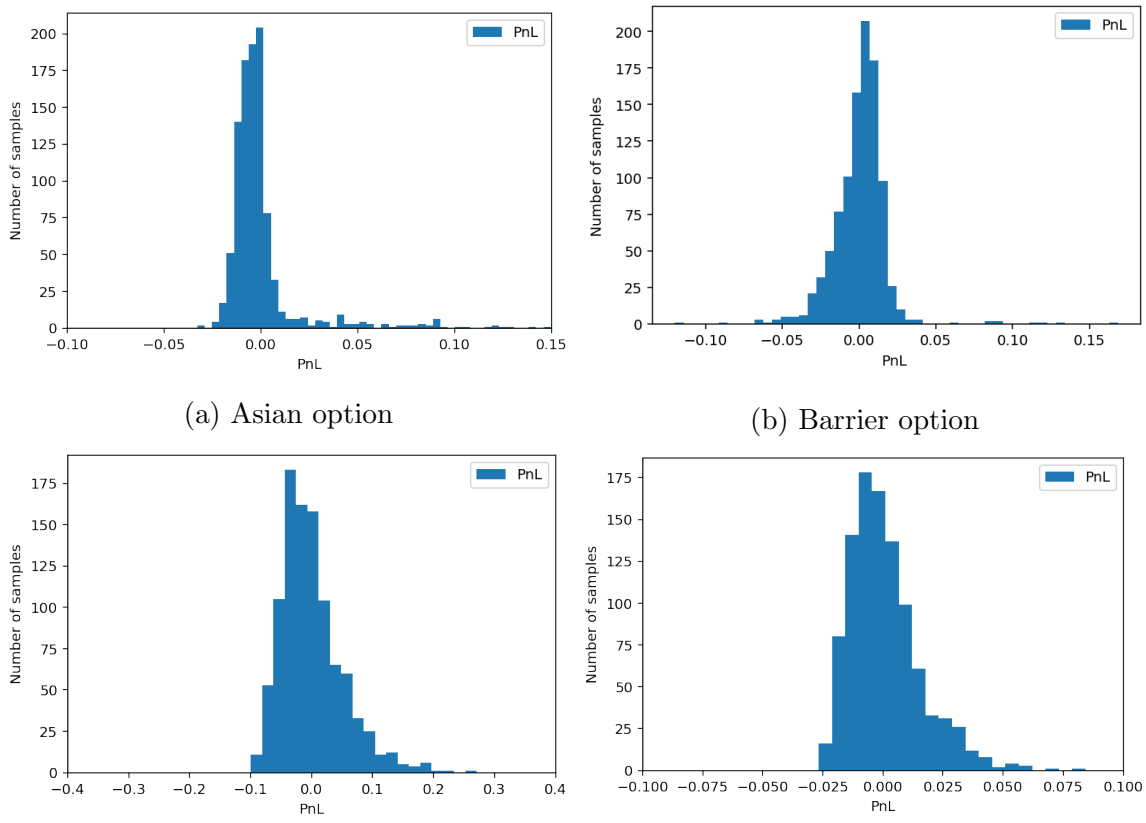

(a) Asian option

(b) Barrier option

(c) Lookback

(d) Variance swap

Figure 3.3: P&L of the hedged portfolio under the Heston model.

3.5.4 Transaction costs

To study the effect of transaction costs, we consider the payoff $F(\widehat{\mathbb{X}}^{LL,<\infty}) = X_T^2$ that was studied in Section 3.5.1. We added fixed quadratic transaction costs (Definition 3.4.5) with $\alpha = 10^{-6}$ and we compared the performance of the signature hedging strategy, obtained by solving (3.5).

Figure 3.5 shows that the P&L of the replicating hedging strategy drops drastically when transaction costs are added, whereas the P&L of the signature hedging strategy is much less affected by these transaction costs. In continuous trading, the P&L of the replicating hedging strategy would incur infinite transaction costs. Moreover, it is known [133, 132] that for small linear transaction costs there is a "no-transaction band" where no trading takes place due to the transaction costs.

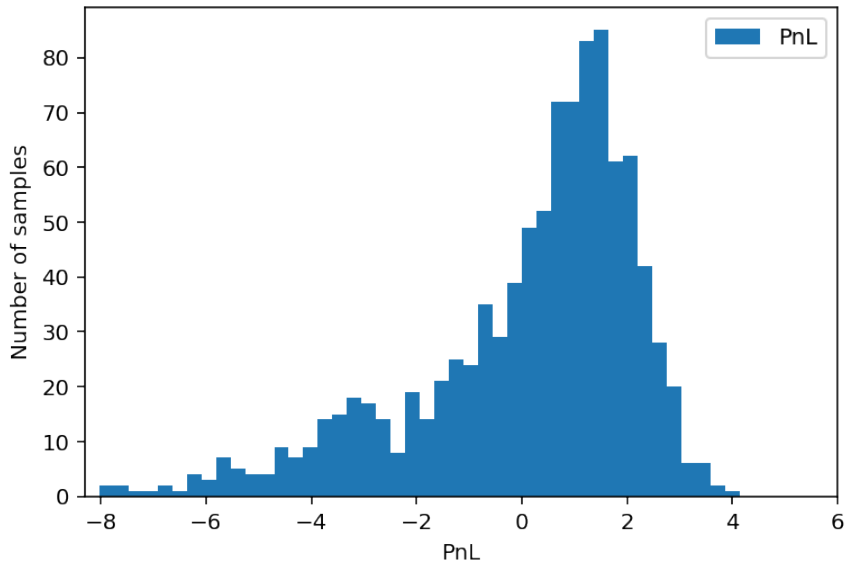

Figure 3.4: P&L of the hedged portfolio for an Asian option, obtained by solving the exponential hedging problem.

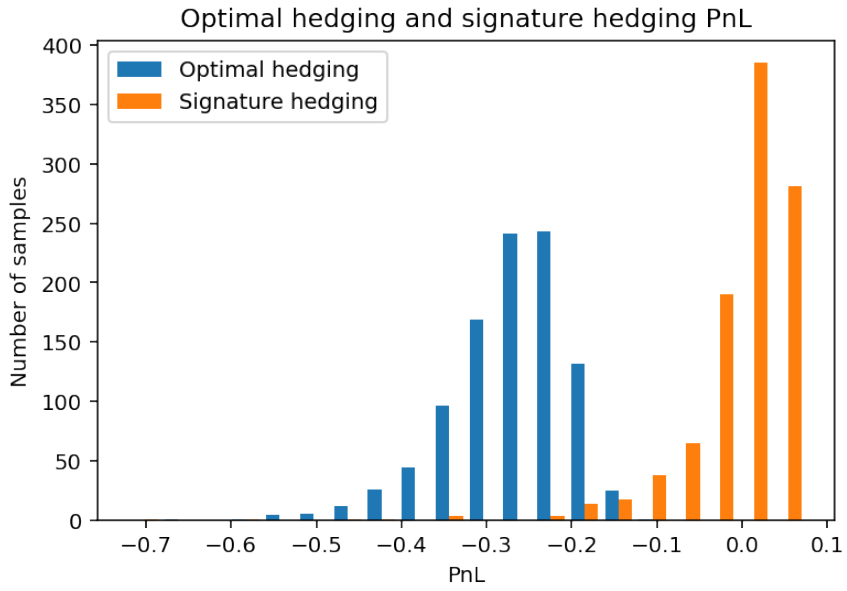

Figure 3.5: P&L of the replicating and signature hedging strategies for $F(\widehat{\mathbb{X}}^{LL,<\infty}) = X_T^2$ with fixed quadratic transaction costs.

3.6 Conclusion

In this chapter we introduce a family of primitive securities called signature payoffs (Definition 3.2.6). In the spirit of Arrow-Debreu, these payoffs approximate arbitrarily well other exotic, path-dependent derivatives. Because signature payoffs are defined as linear combinations of certain iterated integrals, the family of all signature derivatives includes a lot of information about the P&L of all possible dynamic trading strategies.

In Section 3.3, we show that these signature payoffs can be used to reduce the original hard-to-solve optimal hedging problem (PHP) to a polynomial optimisation problem that is numerically easy to solve, (3.1). The only information about the underlying process that is needed to accomplish this is its expected signature – which, in the case where a risk-neutral measure is used, is equivalent to knowing the prices of all signature payoffs. Moreover, our approach is intrinsically nonparametric – we do not need to impose any particular model on the market dynamics.

In the following chapter, we will take the ideas introduced in this chapter a step further and show that provided that the trader has access to prices of a basket of exotic derivatives, she can price and hedge other exotic options in a model-free way.

Chapter 4

Model-free pricing and hedging in discrete time

In Chapter 3 we introduced signature payoffs, a family of primitive financial securities that behave as a *basis* for path-dependent, exotic derivatives. It was mentioned in Chapter 3 that discrete data can be lifted to a lead-lag path. This is precisely the objective of this chapter: to extend the ideas presented in the previous chapter to the discrete time setting. More specifically, we will consider the setting where there exist a finite family of (possibly exotic) payoffs that can be bought at inception ($t = 0$). Then, under minimal assumptions about the market (or in a so-called *model-free* manner) we will show that knowledge of the prices of such payoffs is enough to hedge and estimate prices of other (potentially exotic) payoffs. This chapter is based on the paper [102] published at *Applied Mathematical Finance*, written with Terry Lyons and Sina Nejad.

4.1 Introduction

A framework that has often been considered in model-free finance (see [1, 20, 16, 79]) is that where the market contains a family of financial derivatives written on a risky asset, which can be statically traded. Under certain no-arbitrage assumptions which forbid making profits without taking any risk, one may be interested in studying

whether it is possible to derive a unique price for other financial derivatives. In a model-free framework we do not impose a probability measure on the market dynamics – the objective is to obtain prices of potentially illiquid contracts, without assuming any particular model for the dynamics of the risky asset.

It is well-known ([17]) that if prices of call and put options of all strikes and a fixed maturity are known, one can derive the price of any European contingent claim by writing such contracts as linear combinations of call and put options. Our objective will be to extend Chapter 3 and the idea of approximating complex contracts in terms of simpler, basic contracts to the setting of exotic path-dependent payoffs in discrete time.

Our approach will be analogous to the one exposed in Chapter 3. First, we replicate certain exotic derivatives using a family of primitive securities – signature payoffs. Then, assuming we have access to market prices of a rich enough family of exotic derivatives, we infer the *implied expected signature* – an unobservable quantity that, as we will see, in a sense characterises the entire market dynamics. Finally, we use the inferred implied expected signature to price other exotic payoffs.

Exotic derivatives are typically illiquid and we cannot in general observe the market price of a path-dependent exotic derivative. However, there are multiple data providers that offer consensus market prices of a range of OTC (over-the-counter) exotic derivatives – examples include DeriveXperts Mercure[1] and Totem Markit[2]. These prices reflect the consensus prices from market participants. Moreover, certain FX exotic options such as double no-touch options are liquid and we can observe their market prices. Therefore, exploring the extent to which the market prices of these exotic options contain enough information to price other exotic options is a natural question to ask, and this is precisely the question we address.

[1] `http://www.derivexperts.com/services/mercure`
[2] `https://ihsmarkit.com/products/totem.html`

Throughout this chapter we assume a discrete-time market, where the investor is only allowed to trade at discrete trading times $\mathbb{T} := \{t_i\}_{i=0}^n \subset [0,1]$. Under certain assumptions on the financial market, such markets will have at least a *risk-neutral measure* or *equivalent martingale measure*: a probability measure under which the risky asset is a martingale. However, except in very restrictive settings such risk-neutral measures are not unique. Feasible prices for a financial derivative are given by expectations of the payoff of the contract under risk-neutral measures. Under no-arbitrage assumptions, bounds on feasible prices are also given by the so-called super-hedging prices ([1, Theorem 1.4]). It turns out that under certain assumptions on the market, both pricing approaches coincide – i.e. the supremum of the expected payoff under risk-neutral measures coincides with the infimum super-hedging price ([1, 20, 16]). This is called the pricing-hedging duality.

As we will see in Section 4.3.1, knowledge of the implied expected signature – defined as the expected signature that matches the market prices of all observable financial derivatives – is sufficient to price and hedge financial contracts. The expected signature of a stochastic process ([34, 103, 33, 112]) is an object that plays a similar role to the moments of an $\mathbb{R}^d$-valued random variable. Under certain assumptions, the expected signature determines the law of the stochastic process ([33, 34]), just like under some assumptions the moments of the $\mathbb{R}^d$-valued random variable will determine the law of the random variable. Moreover, knowing the implied expected signature is equivalent to knowing the prices of all signature payoffs. Therefore, the fact that the expected signature determines the law of the market dynamics suggests that the prices of all signature payoffs determine the market dynamics.

The remainder of the chapter is structured as follows. In Section 4.2, we introduce the framework the chapter will be based on, and we state the assumptions we will impose on the market. In Section 4.3 we introduce signature payoffs, which are then

used to price other exotic derivatives. In Section 4.3.1 we explain how these results can be used in practice when one has access to market prices of exotic payoffs. In Section 4.4 we address the problem of hedging exotic derivatives in discrete time, where we use a certain property of signatures (namely, the shuffle product property) to transform a mean-variance optimal hedging problem into a numerically feasible optimisation problem.

4.2 Framework

4.2.1 The market

Let $T > 0$ and let $\mathbb{T} = \{t_i\}_{i=0}^n \subset [0,1]$ with $0 = t_0 < \ldots < t_n = 1$ be the trading times. We assume that the market consists of a single risky asset, although all results generalise to the multi-asset case. The space of (discounted) price paths is defined as $\Omega := \{X : \mathbb{T} \to \mathbb{R}_+ : X_0 = 1\}$. Therefore, we will assume that the initial price is normalised to 1.

Traders make decisions based on past information and the outcome of these decisions is known as new information is revelaed. Thus, we transform price paths to incorporate past and future information. The transformation, known as the *lead-lag transformation* of a path, is studied in [52] and is defined below.

Definition 4.2.1 (Lead-lag transformation)**.** *Given a price path $X \in \Omega$, define its lead-lag transformation $\widehat{X} : [0,1] \to \mathbb{R}^2 \oplus \mathbb{R}$ by the continuous path $\widehat{X}_t = (X_t^b, X_t^f)$, where*

$$\widehat{X}_{2k/2n} := \big((t_k, X_{t_k}), X_{t_k}\big) \in \mathbb{R}^2 \oplus \mathbb{R},$$
$$\widehat{X}_{(2k+1)/2n} := \big((t_k, X_{t_k}), X_{t_{k+1}}\big) \in \mathbb{R}^2 \oplus \mathbb{R},$$

and linear interpolation in between. For each $t \in [0,1]$, we write $\widehat{X}_t = (X_t^b, X_t^f)$ where $X_t^b \in \mathbb{R}^2$ denotes the lag (backward) component and $X_t^f \in \mathbb{R}$ denotes the lead

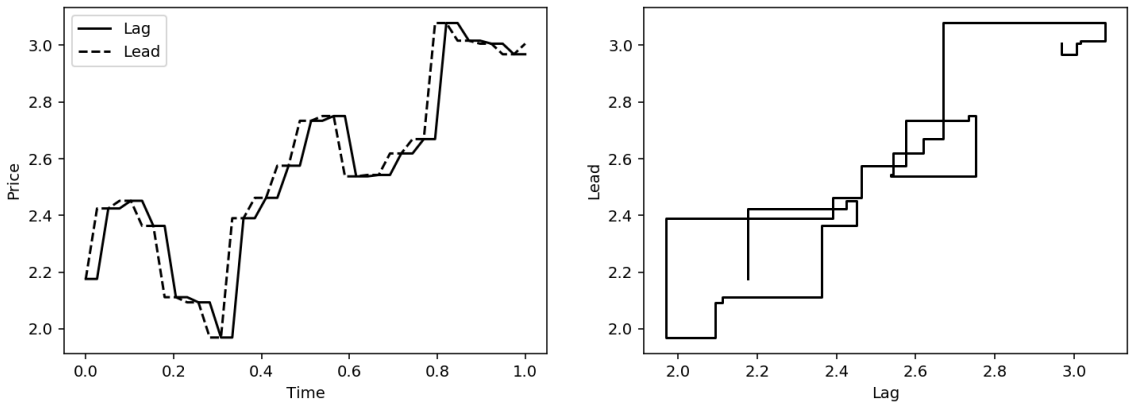

Figure 4.1: Lead-lag transformation of the price path process of a certain financial asset. The figure on the left shows the lead and lag components of the path, and the figure on the right shows the lag component plotted against the lead component.

(forward) component. We define the space of lead-lag price paths $\widehat{\Omega} := \{\widehat{X} : X \in \Omega\}$. *Paths in* $\widehat{\Omega}$ *are continuous and have bounded variation, and hence* $\widehat{\Omega}$ *is a subspace of* $BV([0,1];\mathbb{R}^2 \oplus \mathbb{R})$*, so that* $\widehat{\Omega}$ *is equipped with the norm* $\|\cdot\|_{BV}$*. Similarly, for* $t \in [0,1]$ *we define* $\widehat{\Omega}_t := \{\widehat{X}|_{[0,t]} : \widehat{X} \in \widehat{\Omega}\}$*. We denote by* π^b *and* π^f *the projections* $\pi^b : \mathbb{R}^2 \oplus \mathbb{R} \to \mathbb{R}^2$ *and* $\pi^f : \mathbb{R}^2 \oplus \mathbb{R} \to \mathbb{R}$ *onto the lag and lead components, respectively.*

Figure 4.1 shows the lead-lag transformation of a certain price path. As the name suggests, the lead component (the future) is *leading* the lag component (the past).

Notice that because all paths in $\widehat{\Omega}$ are piecewise linear, $(\widehat{\Omega}, \|\cdot\|_{BV})$ is separable. We consider the measure space $(\widehat{\Omega}, \mathcal{B}(\widehat{\Omega}))$.

It turns out that the signature (Definition 2.1.1) of lead-lag transformed paths uniquely determines the path:

Lemma 4.2.2 (Uniqueness of signature, [75])**.** *The signature map* $\widehat{\Omega} \to T((\mathbb{R}^2 \oplus \mathbb{R}))$ *given by* $\widehat{X} \mapsto \widehat{\mathbb{X}}_{0,1}^{<\infty}$ *is injective.*

4.2.2 Payoffs and no-arbitrage assumptions

Here, we consider exotic, path-dependent payoffs. More specifically, we define payoffs as follows:

Definition 4.2.3 (Payoff)**.** *A payoff is a Borel-measurable function* $G : \widehat{\Omega} \to \mathbb{R}$.

The following are the assumptions we impose on the market.

Assumption 4.2.4. *Let* $\mathcal{M}$ *denote the set of all risk-neutral measures. We assume that the market satisfies the following assumptions:*

1. *There is a finite family of payoffs that can be bought at* $t = 0$, $\mathscr{F} = \{F_i\}_i$.
2. $\mathcal{M} \neq \varnothing$ *(no-arbitrage).*
3. $\mathcal{M}$ *is tight.*
4. $\mathbb{E}^{\mathbb{Q}}[\widehat{\mathbb{X}}_{0,1}^{\leq N}]$ *is well-defined for all* $\mathbb{Q} \in \mathcal{M}$ *and* $N \in \mathbb{N}$.

The assumption that $\mathbb{E}^{\mathbb{Q}}[\widehat{\mathbb{X}}_{0,1}^{\leq N}]$ is well-defined for all $\mathbb{Q} \in \mathcal{M}$ and $N \in \mathbb{N}$ is a mild, nonrestrictive assumption that is considered for technical reasons. It is a path-space analogous of the "moments of all order exist" assumption for finite dimensional random variables.

Definition 4.2.5. *We denote by* $\mathcal{P}$ *the pricing map, given by* $\mathcal{P}(G) := \sup_{\mathbb{Q}\in\mathcal{M}} \mathbb{E}^{\mathbb{Q}}[G(\widehat{X})]$ *for each payoff* G.

Let $F \in \mathscr{F}$. Notice that the mapping $\mathcal{M} \to \mathbb{R}$ given by $\mathbb{Q} \mapsto \mathbb{E}^{\mathbb{Q}}[F(\widehat{X})]$ is constant, so that $\mathcal{P}(F) = \mathbb{E}^{\mathbb{Q}}[F(\widehat{X})]$ for any fixed $\mathbb{Q} \in \mathcal{M}$.

4.2.3 Trading strategies

Trading strategies are investment policies that determine how much of the asset should be held at each time. Moreover, the decision has to be made based on the information about the past only. In our framework, this means that a trading strategy is a real-valued function acting on the lag component of the lead-lag transformation. We make this more precise by introducing the space Λ – see below. A similar space is discussed in [36, 4, 57, 7] and in [48, 11, 120] in the context of finance.

Definition 4.2.6 (Trading strategies)**.** *We denote* $\Lambda := \bigcup_{t\in[0,1]} \pi^b(\widehat{\Omega}_t)$*, where* π^b *is the projection onto the lag component defined in Definition 4.2.1. This is a metric space with a certain metric. The space of trading strategies is defined as* $\mathcal{T} := C(\Lambda; \mathbb{R})$*.*

Intuitively, the space $\mathcal{T}$ consists of all non-anticipative processes; trading should only be done with past information, because no information about the future can be used.

Remark 4.2.7. The space Λ is a Polish space (see [57]).

The outcome of a trading strategy $\theta \in \mathcal{T}$ depends on the trajectory of the price path $\widehat{X} \in \widehat{\Omega}$. The gains of the trading strategy θ are defined below.

Definition 4.2.8. *For* $\theta \in \mathcal{T}$ *and* $X \in \widehat{\Omega}$*, the gains of the trading strategy* θ *are given by*

$$(\theta \bullet \widehat{X})_{\mathbb{T}} := \sum_k \theta(X^b|_{[0,t_k]})(X_{t_{k+1}} - X_{t_k}).$$

In this chapter we allow semi-static hedging. In other words, the trader is allowed to buy or sell the contracts in $\mathscr{F}$, which can only be bought or sold at $t = 0$, as well as to dynamically trade the underlying asset. The performance of the hedging strategy will then be determined by three factors: the price that is paid for each of the payoffs in $\mathscr{F}$, the gains or losses that stem from holding the payoffs in $\mathscr{F}$, and the gains that come from dynamically trading the underlying asset.

Definition 4.2.9. *A hedging strategy is a pair* $h = ((a_F)_{F\in\mathscr{F}}, \theta)$*, where* $a_F \in \mathbb{R}$ *for all* $F \in \mathscr{F}$ *and* $\theta \in \mathcal{T}$*. The Profit and Loss (P& L) of* h *is defined by*

$$\begin{aligned} V_h : \widehat{\Omega} &\to \mathbb{R} \\ \widehat{X} &\mapsto \sum_{F\in\mathscr{F}} a_F(F(\widehat{X}) - \mathcal{P}(F)) + (\theta \bullet X)_{\mathbb{T}}. \end{aligned}$$

We denote by $\mathcal{H}$ the space of hedging strategies. A payoff $G : \widehat{\Omega} \to \mathbb{R}$ is said to be attainable if there exists a hedging strategy $h \in \mathcal{H}$ such that $G \equiv V_h$ on $\widehat{\Omega}$. We denote the space of all attainable payoffs by $\mathcal{A}$. Similarly, we define, for $\varepsilon > 0$, $\mathcal{A}_\varepsilon := \{G : \widehat{\Omega} \to \mathbb{R}$ payoff $\mid \exists H \in \mathcal{A}$ such that $\|G - H\|_{L^\infty(\widehat{\Omega})} < \varepsilon\}$ the space of almost attainable payoffs.

4.3 Pricing

We now define a special class of continuous payoffs – the family of signature payoffs. These were first introduced in [116, 101], and as Arrow-Debreu securities [5, 42] they are *primitive securities* in the sense that path-dependent exotic derivatives can be approximated by these signature payoffs.

Definition 4.3.1 (Signature payoff)**.** *Let $\ell \in T((\mathbb{R}^2 \oplus \mathbb{R})^*)$. The signature payoff $\mathcal{S}^\ell$ is defined by*

$$\begin{aligned} \mathcal{S}^\ell : \widehat{\Omega} &\to \mathbb{R} \\ \widehat{X} &\mapsto \langle \ell, \widehat{\mathbb{X}}_{0,1}^{<\infty} \rangle. \end{aligned}$$

In other words, a signature payoff is a financial derivative whose payoff is given as a linear combination of certain iterated integrals. Because we are using iterated integrals against the price path, the family of all signature payoffs effectively contains all possible dynamic hedging strategies. Therefore, it is not surprising that the space of signature payoffs is big, as a subspace of all payoffs.

The following Proposition shows that signature payoffs are primitive securities. Indeed, the proposition establishes that continuous payoffs can be locally approximated by signature payoffs:

Proposition 4.3.2. *Let $G : \widehat{\Omega} \to \mathbb{R}$ be a continuous payoff, and let $\mathcal{K} \subset \widehat{\Omega}$ be compact. Given any $\varepsilon > 0$, there exists a signature payoff $\mathcal{S}^\ell$ with $\ell \in T((\mathbb{R}^2 \oplus \mathbb{R})^*)$ such that*

$$|G(\widehat{X}) - \mathcal{S}^\ell(\widehat{X})| < \varepsilon \quad \forall\, \widehat{X} \in \mathcal{K}.$$

We have the following immediate corollary.

Corollary 4.3.3. *Let $G : \widehat{\Omega} \to \mathbb{R}$ be a continuous payoff. Let $\varepsilon > 0$. Then, there exists a compact $\mathcal{K}_\varepsilon \subset \widehat{\Omega}$ and a signature payoff $\mathcal{S}^\ell$ with $\ell \in T((\mathbb{R}^2 \oplus \mathbb{R})^*)$ such that*

1. *$\mathbb{Q}(\widehat{X} \in \mathcal{K}_\varepsilon) > 1 - \varepsilon$ for all risk-neutral measures $\mathbb{Q} \in \mathcal{M}$,*
2. *$|G(\widehat{X}) - \mathcal{S}^\ell(\widehat{X})| < \varepsilon$ for all $\widehat{X} \in \mathcal{K}_\varepsilon$.*

The corollary states that there exists a *large* compact set – large in the sense that with very high probability, the price path will belong to the compact set – such that on the compact set the exotic derivative is very close to a signature payoff.

Moreover, prices of payoffs can be well-approximated using signature payoffs:

Proposition 4.3.4. *Set $\varepsilon > 0$. Let $G : \widehat{\Omega} \to \mathbb{R}$ be a $\mathbb{Q}$-integrable payoff for all $\mathbb{Q} \in \mathcal{M}$. Assume G is either in $\mathcal{A}_{\varepsilon/4}$ (i.e. almost attainable) or bounded. Then, there exists a compact $\mathcal{K}_\varepsilon \subset \widehat{\Omega}$ and a signature payoff $\mathcal{S}^\ell$ such that, for $\mathcal{L} := \mathbb{1}_{\mathcal{K}_\varepsilon}\mathcal{S}^\ell : \widehat{\Omega} \to \mathbb{R}$,*

1. *$\mathbb{Q}(\widehat{X} \in \mathcal{K}_\varepsilon) > 1 - \varepsilon$ for all $\mathbb{Q} \in \mathcal{M}$.*
2. *$\mathcal{P}(|G - \mathcal{L}|) = \sup_{\mathbb{Q}\in\mathcal{M}} \mathbb{E}^{\mathbb{Q}}[|G(\widehat{X}) - \mathcal{L}(\widehat{X})|] < \varepsilon$.*

4.3.1 The implied expected signature

Let $\varepsilon > 0$. By Proposition 4.3.4, we may pick a compact set $\mathcal{K}_\varepsilon \subset \widehat{\Omega}$ and a family of signature payoffs $\mathscr{L} = \{\mathcal{S}^{\ell_F}\}_{F\in\mathscr{F}}$ with $\ell_F \in T((\mathbb{R}^2 \oplus \mathbb{R})^*)$, such that:

1. $\mathbb{Q}(\widehat{X} \in \mathcal{K}_\varepsilon) > 1 - \varepsilon$ for all $\mathbb{Q} \in \mathcal{M}$,

2. $\mathcal{P}(|F - \mathbb{1}_{\mathcal{K}_\varepsilon}\mathcal{S}^{\ell_F}|) < \varepsilon$ for all $F \in \mathscr{F}$.

Fix a risk-neutral measure $\mathbb{Q} \in \mathcal{M}$ now. We then have:

$$\langle \ell_F, \mathbb{E}^{\mathbb{Q}}[\widehat{\mathbb{X}}_{0,1}^{<\infty} \,;\, \mathcal{K}_\varepsilon] \rangle \approx \mathcal{P}(F) \quad \forall F \in \mathscr{F}. \tag{4.1}$$

Choose $N \in \mathbb{N}$ sufficiently large so that $\ell_F \in T^{(N)}(\mathbb{R}^2 \oplus \mathbb{R})$ for all $F \in \mathscr{F}$, which is possible because $\mathscr{F}$ is finite and by definition of tensor algebra (Definition 1.1.2). Then, (4.1) becomes:

$$\langle \ell_F, \mathbb{E}^{\mathbb{Q}}[\widehat{\mathbb{X}}_{0,1}^{\leq N} \,;\, \mathcal{K}_\varepsilon] \rangle \approx \mathcal{P}(F) \quad \forall F \in \mathscr{F}. \tag{4.2}$$

If the risk-neutral measures are unknown and we only observe market prices for the exotic payoffs $\mathscr{F}$, we may look for the expected signature that best matches the observed prices – i.e. the best expected signature fit for the approximation (4.2).

To what extent this can be achieved depends on the invertibility of the map

$$T^{(N)}(\mathbb{R}^2 \oplus \mathbb{R}) \to \mathbb{R}^{|\mathscr{F}|} \tag{4.3}$$

$$\mathbf{a} \mapsto (\langle \ell_F, \mathbf{a} \rangle)_{F \in \mathscr{F}}, \tag{4.4}$$

as given a vector of prices $(\mathcal{P}(F))_{F \in \mathscr{F}} \in \mathbb{R}^{|\mathscr{F}|}$ we are looking for $\mathbf{a} \in T^{(N)}(\mathbb{R}^2 \oplus \mathbb{R})$ such that $\langle \ell_F, \mathbf{a} \rangle \approx \mathcal{P}(F)$ for all $F \in \mathscr{F}$. If the family of exotic payoffs $\mathscr{F}$ is rich enough (in the sense that the linear map (4.3) has high rank) we may be able to estimate the expected signature that best matches the observed prices for such exotic payoffs. We call this expected signature the *implied expected signature*, in analogy with the *implied volatility*.

If we are given a (potentially illiquid) payoff $G : \widehat{\Omega} \to \mathbb{R}$ that is close to the space of attainable payoffs $\mathcal{A}$ (Definition 4.2.9) and that satisfies the assumptions of Proposition 4.3.4, we may estimate its price by $\mathcal{P}(G) \approx \langle \ell_G, \mathbf{E} \rangle$ where ℓ_G is given by

Proposition 4.3.4 and $\mathbf{E}$ is the implied expected signature that was estimated from market data.

This approach provides a method to estimate prices of certain exotic payoffs from observed prices of other exotics in a model-free manner. The effectiveness and accuracy of the method depends on the richness of the payoffs $\mathscr{F}$ with observable prices.

4.4 Hedging

When one buys or sells a financial derivative one is immediately exposed to an uncertain future cash flow. Sometimes this exposure is desired, such as when the intention is to speculate on the underlying asset. However, sometimes this exposure is undesired and the trader would like to hedge that risk away.

If the market is complete, it is possible to exactly hedge the derivative by trading on the underlying. However, completeness is a strong assumption and markets are not, in general, complete. In this case, the trader may try to find a hedging strategy that is optimal in the sense that it minimises a certain cost function. This cost function would be chosen by the trader, depending on her risk preferences.

Let $F \in \mathscr{F}$, and fix $\mathbb{Q} \in \mathcal{M}$. Assume that $F \in L^2(\mathbb{Q})$. In this section, we will study an optimal hedging problem, known as the mean-variance problem:

$$\inf_{\theta \in \mathcal{T}} \mathbb{E}^{\mathbb{Q}} \left[\left(F(\widehat{X}) - p_0 - (\theta \bullet \widehat{X})_{\mathbb{T}} \right)^2 \right]. \tag{4.5}$$

where $p_0 \in \mathbb{R}$ is the initial cash, typically the price of F. We will do so by extending the ideas presented in Chapter 3 to the discrete time setting.

4.4.1 A signature-version of the optimal hedging problem

In this section, we will reduce the original problem (4.5) to a signature version of the problem which will turn out to be numerically easier to solve. Following the idea of

approximating functions on paths by linear functions on the signature that was used in Section 4.3, we will, for any given $\varepsilon > 0$:

1. Pick a compact set $\mathcal{K}_\varepsilon$ such that $\mathbb{Q}(\widehat{X} \in \mathcal{K}_\varepsilon) > 1 - \varepsilon$ and $\mathbb{E}^{\mathbb{Q}}[|F(\widehat{X})|^2 \; ; \; \mathcal{K}_\varepsilon^c] < \varepsilon$.

2. Replace the payoff F with a signature payoff $\mathcal{S}^{\ell_F}$ that approximates F on $\mathcal{K}_\varepsilon$.

3. Replace the trading strategy $\theta \in \mathcal{T}$ with a *signature trading strategy* that approximates θ on $\mathcal{K}_\varepsilon$, given by $X^b|_{[0,t]} \mapsto \langle \ell, \mathbb{X}_{0,t}^{b,\leq N} \rangle$, where $\ell \in T((\mathbb{R}^2)^*)$ and $\mathbb{X}^{b,\leq N}$ is the signature of the lag component X^b.

Step 2 is justified by Proposition 4.3.4. Similarly, Step 3 is justified by the following proposition, which is the discrete-time version of Proposition 3.3.8 from Chapter 3:

Proposition 4.4.1. *Let $\varepsilon > 0$, and let $\theta \in \mathcal{T}$ be a trading strategy. Then, there exists a compact set $\mathcal{K}_\varepsilon \subset \widehat{\Omega}$, independent of θ, and a signature trading strategy given by $X^b|_{[0,t]} \mapsto \langle \ell, \mathbb{X}_{0,t}^{b,<\infty} \rangle$ with $\ell \in T((\mathbb{R}^2)^*)$, such that:*

1. *$\mathbb{Q}(\widehat{X} \in \mathcal{K}_\varepsilon) > 1 - \varepsilon$ for all $\mathbb{Q} \in \mathcal{M}$,*

2. *$|\theta(X^b|_{[0,t]}) - \langle \ell, \mathbb{X}_{0,t}^{b,<\infty} \rangle| < \varepsilon$ for all $\widehat{X} \in \mathcal{K}_\varepsilon$ and $t \in [0,1]$.*

Then, we may estimate the solution of the optimal hedging problem (4.5) by the following *optimal signature hedging problem*:

$$\inf_{\ell \in T((\mathbb{R}^2)^*)} \mathbb{E}^{\mathbb{Q}} \left[\left(\langle \ell_F, \widehat{\mathbb{X}}_{0,1}^{<\infty} \rangle - p_0 - \sum_k \langle \ell, \mathbb{X}_{0,t_k}^{b,\infty} \rangle (X_{t_{k+1}} - X_{t_k}) \right)^2 \; ; \; \mathcal{K}_\varepsilon \right]. \tag{4.6}$$

4.4.2 Solving the optimal signature hedging problem

We will begin by rewriting (4.6) appropriately. We will do so by using the following lemma, according to which the gains of a dynamic trading strategy $\theta \in \mathcal{T}$, which is

defined by a discrete integral, can be written exactly as an integral against the lead component. This was observed in [52], and justifies the introduction of the lead-lag transformation.

Lemma 4.4.2. *Let $\theta \in \mathcal{T}$. Then, $(\theta \bullet \widehat{X})_{\mathbb{T}} = \int_0^1 \theta(X^b|_{[0,t]})dX_t^f$ for all $\widehat{X} \in \widehat{\Omega}$.*

Proof. We have

$$\begin{aligned}\sum_{t_k} \theta(X^b|_{[0,t_k]})(X_{t_{k+1}} - X_{t_k}) = \sum_k \int_{\frac{2k}{2n}}^{\frac{2k+1}{2n}} \theta(X^b|_{[0,t]})dX_t^f = \sum_k \int_{\frac{2k}{2n}}^{\frac{2k+2}{2n}} \theta(X^b|_{[0,t]})dX_t^f \\ = \int_0^1 \theta(X^b|_{[0,t]})dX_t^f.\end{aligned}$$

□

By the previous lemma, (4.6) can be rewritten as:

$$\inf_{\theta\in\mathcal{T}} \mathbb{E}^{\mathbb{Q}}\left[\left(F(\widehat{X}) - p_0 - \int_0^1 \langle \ell, \mathbb{X}_{0,t}^{b,<\infty}\rangle dX_t^f\right)^2 \;;\; \mathcal{K}_\varepsilon\right]. \tag{4.6'}$$

The objective will be to solve (4.6') using signature payoffs and the implied expected signature.

Using the shuffle product property, the optimisation problem (4.6') can then be rewritten as:

$$\begin{aligned}&\inf_{\ell\in T((\mathbb{R}^2)^*)} \mathbb{E}^{\mathbb{Q}}\left[\left(\langle \ell_F, \widehat{\mathbb{X}}_{0,1}^{<\infty}\rangle - p_0 - \int_0^1 \langle \ell, \mathbb{X}_{0,t}^{b,<\infty}\rangle dX_t^f\right)^2 \;;\; \mathcal{K}_\varepsilon\right] \\ &= \inf_{\ell\in T((\mathbb{R}^2)^*)} \mathbb{E}^{\mathbb{Q}}\left[\left(\langle \ell_F, \widehat{\mathbb{X}}_{0,1}^{<\infty}\rangle - p_0 - \langle \ell\mathbf{3}, \widehat{\mathbb{X}}_{0,1}^{<\infty}\rangle\right)^2 \;;\; \mathcal{K}_\varepsilon\right] \\ &= \inf_{\ell\in T((\mathbb{R}^2)^*)} \mathbb{E}^{\mathbb{Q}}\left[\langle (\ell_F - p_0\varnothing - \ell\mathbf{3})^{ш2}, \widehat{\mathbb{X}}_{0,1}^{<\infty}\rangle \;;\; \mathcal{K}_\varepsilon\right] \\ &= \inf_{\ell\in T((\mathbb{R}^2)^*)} \left\langle (\ell_F - p_0\varnothing - \ell\mathbf{3})^{ш2}, \mathbb{E}^{\mathbb{Q}}\left[\widehat{\mathbb{X}}_{0,1}^{<\infty} \;;\; \mathcal{K}_\varepsilon\right]\right\rangle.\end{aligned}$$

Therefore, we have proved the following:

Proposition 4.4.3. *Let $\ell_F \in T((\mathbb{R}^2 \oplus \mathbb{R})^*)$ and consider the corresponding signature payoff $\mathcal{S}^{\ell_F}$. Fix $\mathbb{Q} \in \mathcal{M}$, $\mathcal{K} \subset \widehat{\Omega}$ compact and let $p_0 \in \mathbb{R}$ be the initial cash. Then, the solution of the mean-variance problem* (4.6) *is given by the solution of the following optimisation problem:*

$$\inf_{\ell \in T((\mathbb{R}^2)^*)} \left\langle (\ell_F - p_0 \varnothing - \ell \mathbf{3})^{\text{ш}2}, \mathbb{E}^{\mathbb{Q}} \left[\widehat{\mathbb{X}}_{0,1}^{<\infty} \ ; \ \mathcal{K} \right] \right\rangle. \tag{4.7}$$

The original problem (4.6) nonlinearly involves the control, ℓ, and the stochastic component, X_t. Proposition 4.4.3 separates this nonlinear dependency by rewriting the original optimisation problem as a deterministic linear functional, that depends on the control but not on the stochastic component of the problem, applied to the expected signature of the price path, that depends on the stochastic component and not on the control.

In practice, for obvious computational reasons, one would consider the truncated version of the optimisation problem:

$$\inf_{\ell \in T^{\lfloor N/2 \rfloor}((\mathbb{R}^2)^*)} \left\langle (\ell_F - p_0 \varnothing - \ell \mathbf{3})^{\text{ш}2}, \mathbb{E}^{\mathbb{Q}} \left[\widehat{\mathbb{X}}_{0,1}^{\leq N} \right] \right\rangle, \quad \text{for } N \in \mathbb{N}. \tag{4.8}$$

The optimisation problem (4.8) is a finite-dimensional optimisation problem. In fact, it consists of finding the global minimum of a quadratic polynomial, a task that is fast to solve computationally. Therefore, the original hard-to-solve optimal hedging problem (4.5) has been reduced to a more tractable, computationally feasible optimisation problem (4.8). Moreover, we can use the *implied* expected signature that was introduced in Section 4.3.1 as an estimate of $\mathbb{E}^{\mathbb{Q}} \left[\widehat{\mathbb{X}}_{0,1}^{\leq N} \right]$. This allows us to solve the optimal hedging problem in a model-free way, from market data.

Remark 4.4.4. This approach generalises for *almost attainble* payoffs $G \in \mathcal{A}_\varepsilon$.

4.5 Numerical experiment

4.5.1 Approximating payoffs by signature payoffs

In this section we discuss how to approximate a payoff $G : \widehat{\Omega} \to \mathbb{R}$ by a signature payoff $\mathcal{S}^\ell : \widehat{\Omega} \to \mathbb{R}$.

We begin by choosing a dataset of paths $\{\widehat{X}^{(i)}\}_{i=1}^n \subset \widehat{\Omega}$. In the experiments in this section, we generated $n = 10,000$ paths from a Black–Scholes model with a constant volatility ranging from 5% to 40%, terminal time 1 year and daily sampling – as we will see, the approximating signature payoff generalised well even to non-Black–Scholes paths.

Once the set of paths is chosen, we compute the signatures $\{(\widehat{\mathbb{X}}_{0,1}^{\leq N})^{(i)}\}_{i=1}^n$ of order $N \in \mathbb{N}$. In the following experiments we fixed $N = 5$. There are multiple publicly available, open source libraries to compute signatures, such as `esig`[3] or `iisignature`[4].

Finally, one can apply a linear regression of the signatures $\{(\widehat{\mathbb{X}}_{0,1}^{\leq N})^{(i)}\}_{i=1}^n$ against the payoff cash flows $\{G(\widehat{X}^{(i)})\}_{i=1}^n$ to find a linear functional $\ell \in T^{(N)}(\mathbb{R}^2 \oplus \mathbb{R})$ such that

$$\langle \ell, (\widehat{\mathbb{X}}_{0,1}^{\leq N})^{(i)} \rangle \approx G(\widehat{X}^{(i)}) \quad \forall\, i = 1, \ldots, n.$$

4.5.2 Pricing from market data

In this section we implement the proposed approach to estimate the implied expected signature, which we employ to price exotic options. The family $\mathfrak{F}$ of exotic options available to buy at $t = 0$ (see Assumption 4.2.4) that was considered consists of 100 exotic and vanilla payoffs with maturity 1 year. More specifically, we consider European call options, barrier up-and-out call options, barrier up-and-in call options and variance options (25 payoffs of each type). The size of the dataset, as well as

[3] `https://pypi.org/project/esig/`
[4] `https://github.com/bottler/iisignature`

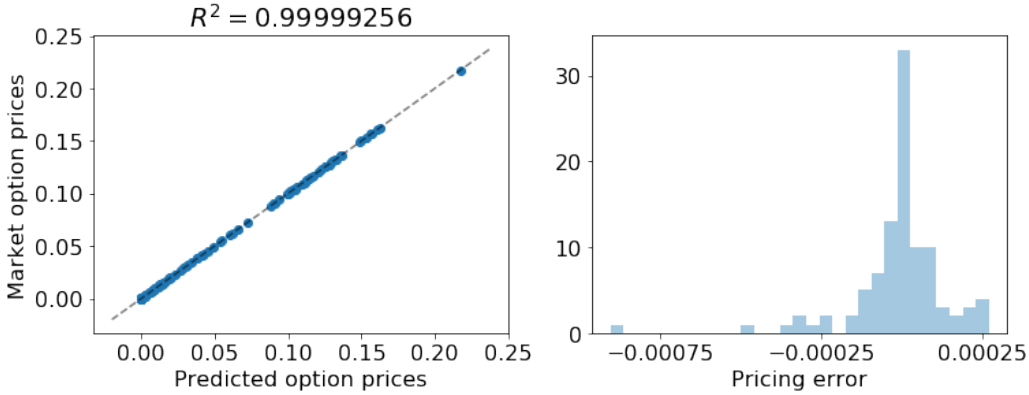

(a) Hull–White model.

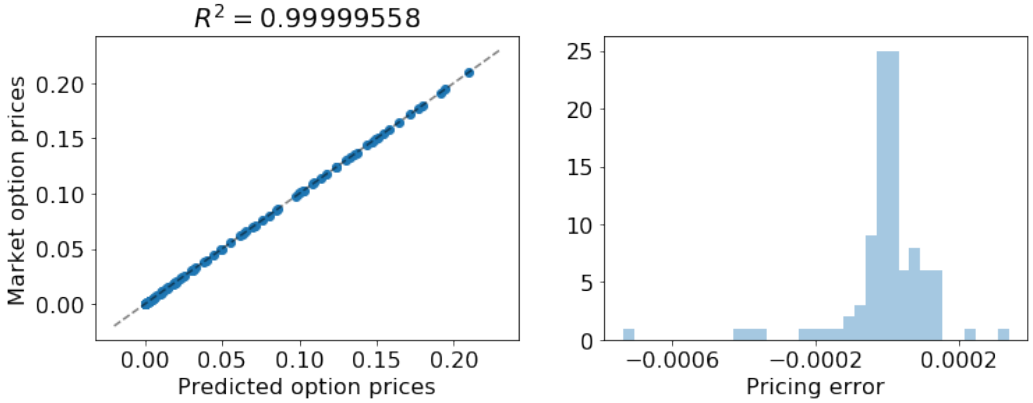

(b) GARCH model.

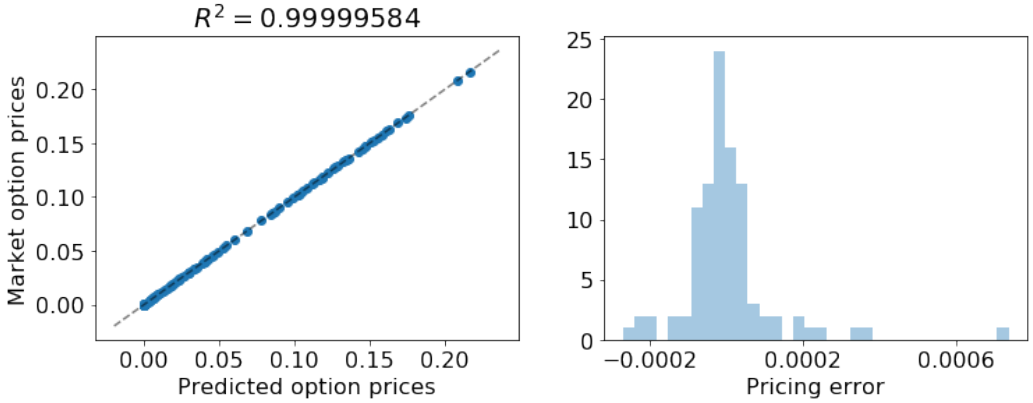

(c) Rough volatility model.

Figure 4.2: Accuracy of the estimated market prices for different market dynamics. The scatter plot on the left shows the predicted option price (horizontal axis) against the actual market price (vertical axis). The histogram on the right shows the pricing error.

the payoff types, is similar to that offered by consensus market price providers. The discrete timeline $\mathbb{T}$ we consider is a timeline consisting of all trading days in that year.

To generate prices of $\mathfrak{F}$, we considered prices generated by three stochastic volatility models: Hull–White model [81], GARCH model [9, 8] and the rough volatility model [59]. In all cases the market dynamics are unknown to the trader because the trader has access to the prices of exotics in $\mathfrak{F}$ but has no knowledge about the market dynamics that generated those prices. Therefore, the only information the trader may use is the market prices of $\mathfrak{F}$.

We used the procedure described in Section 4.5.1 to approximate each payoff in $\mathfrak{F}$ by a signature payoff, and then we followed the numerical method proposed in Section 4.3.1 to estimate the implied expected signature. Signatures of order $N = 5$ are considered.

We generated prices of 100 new exotic derivatives (different from those in $\mathfrak{F}$) using each of the three models (i.e. Hull–White model, GARCH and rough volatility). We used the implied expected signature to predict the prices of these 100 new exotic payoffs, for which we obtain accurate results for all models – see Figure 4.2. The R^2 of the predictions is very high for the three models.

The effectiveness of the methodology depends on how rich the class of exotic payoffs $\mathfrak{F}$ is, as discussed in Section 4.3.1. For instance, the larger the family of exotics $\mathfrak{F}$ is, the more accurate the estimation of the implied expected signature is, which leads to more accurate market price predictions. This is shown on the first figure of Figure 4.3, where the R^2 of the predicted market prices of exotic payoffs is shown as a function of the size of the available exotic payoffs $\mathfrak{F}$. As we see, the performance is not very good when the size of $\mathfrak{F}$ is small, but it rapidly improves as the size of $\mathfrak{F}$ increases. On the other hand, the types of exotic payoffs that are considered also matters. For instance, if $\mathfrak{F}$ only includes vanilla options we have information

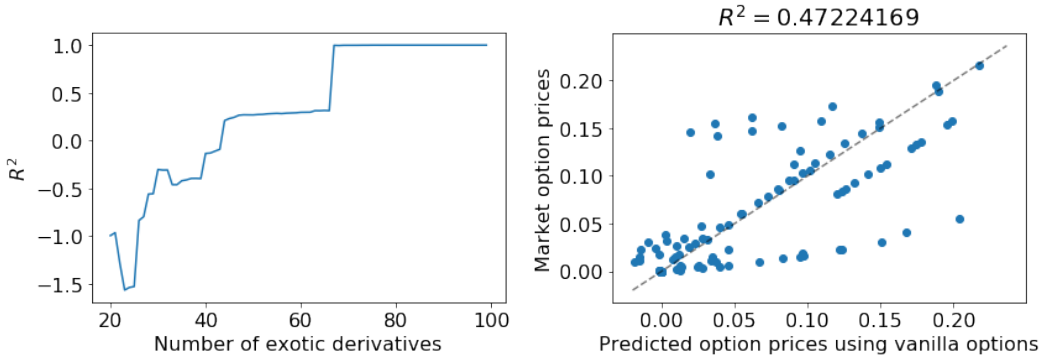

Figure 4.3: On the left, R^2 of the predicted market prices of the exotic functions, as a function of the size of available market prices $|\mathfrak{F}|$. On the right, the predicted market prices of exotic options (horizontal axis) against the actual market price (vertical axis) when $\mathfrak{F}$ is only formed by vanilla options. The R^2 is very poor when only vanilla options are used. Both plots were obtained for prices computed under the Hull–White model.

	Vanillas only in $\mathfrak{F}$	**Vanilla and exotic options in $\mathfrak{F}$**		
	$\lvert\mathfrak{F}\rvert = 100$	$\lvert\mathfrak{F}\rvert = 33$	$\lvert\mathfrak{F}\rvert = 66$	$\lvert\mathfrak{F}\rvert = 100$
R^2 of predictions	0.472242	0.461247	0.313326	**0.999993**
Mean-squared error	1.88×10^{-3}	1.81×10^{-2}	1.54×10^{-6}	$\mathbf{1.48 \times 10^{-8}}$
Mean absolute error	2.94×10^{-2}	7.89×10^{-2}	6.04×10^{-4}	$\mathbf{3.39 \times 10^{-5}}$

Table 4.1: Comparison between different choices of available payoffs $\mathfrak{F}$. In the first column only 100 vanilla options are in $\mathfrak{F}$ and in the second column, $\mathfrak{F}$ consists of both vanilla and exotic options (barrier up-and-out call options, barrier up-and-in call options and variance option). The table suggests that a large family of payoffs $\mathfrak{F}$ that includes exotic options leads to the best performance.

about the marginal distribution of the price path, but it will not, in general, have information about the transition probabilities of prices. Therefore, vanilla options are not sufficient to price other exotic options. This is illustrated on the second figure of Figure 4.3, where the accuracy of predicted market prices of exotic payoffs is shown in the case where $\mathfrak{F}$ only contains vanilla options.

Table 4.1 includes the accuracy of the predicted prices, measured with different metrics, for different choices of $\mathfrak{F}$. As the table indicates, a large family of payoffs $\mathfrak{F}$ that includes vanilla and exotic options leads to the best performance.

4.5.3 Hedging from market data

In this section we implement the methodology proposed in Section 4.4 to solve the optimal hedging problem (4.5). As in Section 4.5.2, the family $\mathscr{F}$ of exotic options available to buy at $t = 0$ (see Assumption 4.2.4) that was considered consists of 25 barrier up-and-out call options, 25 barrier up-and-in call options, 25 variance options and 25 European call options (100 payoffs in total).

We generated prices of $\mathscr{F}$ using a Heston model [77], which is a stochastic volatility model. As in Section 4.5.2, the market dynamics are unknown to the trader as the only information available to the trader is the prices of exotics in $\mathscr{F}$. Following Section 4.3.1, we use this information to estimate the implied expected signature, which is used to solve the optimal hedging problem (4.5). Figure 4.4 shows the P&L of the hedged portfolio corresponding to various payoffs.

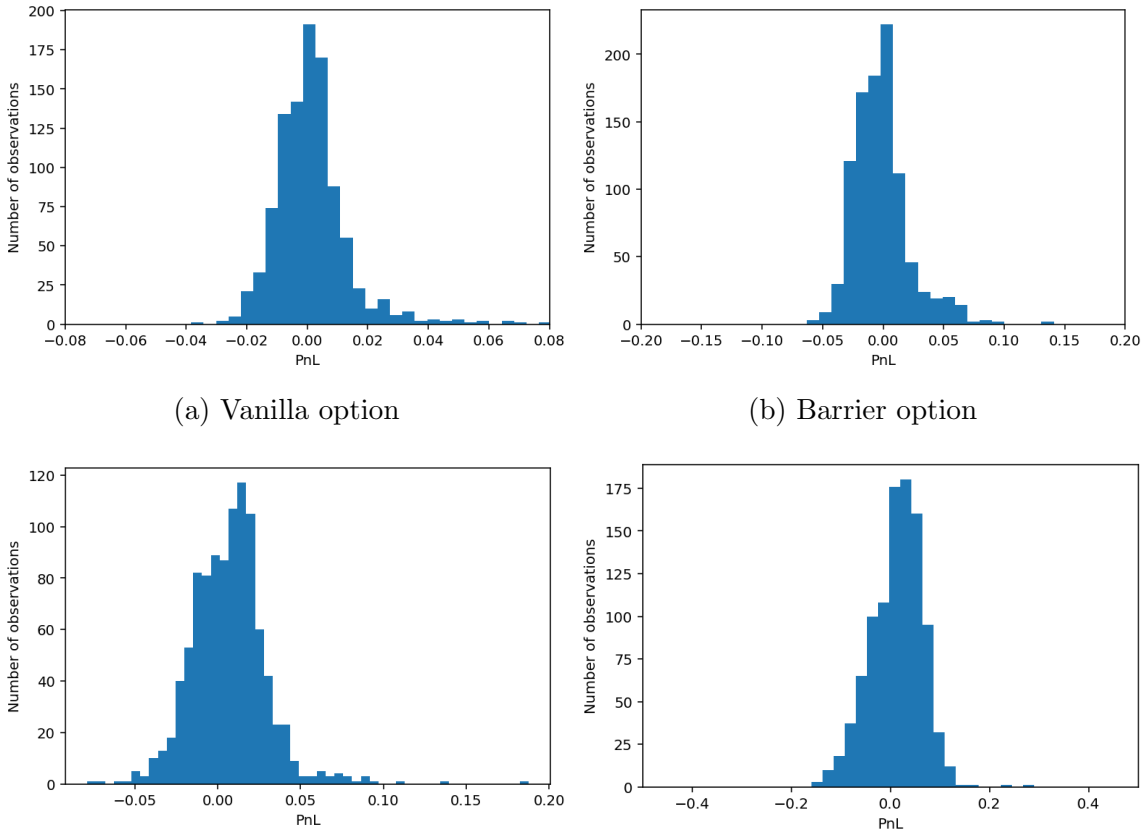

(a) Vanilla option

(b) Barrier option

(c) Asian option

(d) Variance swap

Figure 4.4: P&L of the hedged portfolio of various payoffs, obtained using the implied expected signature.

Chapter 5

Optimal execution

The last two chapters showed how signatures can be used to numerically price and hedge financial derivatives. This chapter, together with the next one, will discuss applications of signatures in the problem of optimal execution, which has gained popularity in the last few years. This chapter is derived from the paper [86], which was accepted at *SIAM Journal on Financial Mathematics*, written in collaboration with Terry Lyons and Jasdeep Kalsi.

We present a method for obtaining approximate solutions to the problem of optimal execution, based on a signature method. The framework is general, only requiring that the price process is a geometric rough path and the price impact function is a continuous function of the trading speed. Following an approximation of the optimisation problem, we calculate an optimal solution for the trading speed in the space of linear functions on a truncation of the signature of the price process. We provide strong numerical evidence illustrating the accuracy and flexibility of the approach. Our numerical investigation both examines cases where exact solutions are known, demonstrating that the method accurately approximates these solutions, and models where exact solutions are not known. In the latter case, we obtain favourable comparisons with standard execution strategies.

5.1 Introduction

5.1.1 Overview

The problem of optimal execution has attracted much interest following the original work on the problem by Bertsimas and Lo in [12] and Almgren and Chriss in [3]. The aim is to model how one should send orders to the market in order to transition from holding one portfolio to another. Typically the case where an investor simply wishes to acquire/liquidate shares in a single asset is considered. There are two competing factors to the optimisation. Firstly, the investor has pressure to trade quickly. Trading more at later times would mean accepting more risk, as the future prices are uncertain. On the other hand, trading evenly across time also has its benefits due to the nature of market price impact. The investor should consider the liquidity at the desirable prices – placing a large order now could result in "walking the book" and accepting unfavourable prices for a large portion of their trade.

The key features in any optimal execution model are the dynamics of the price process at which the trader can execute her trades, P_t, and some definition of a good strategy for the trader. The process P_t is a function of the history of the trading speed until time t, together with some additional driving processes. Typically, we have that P_t is given by the sum of an underlying price process added to a price impact function. The price impact function depends on the history of the investor's trading speed, and it determines how much the price at which the trader can execute has changed as a consequence of that. Classical choices of price impact functions include *temporary* versions, which depend only on the speed at which the trader wishes to trade at that time and try to model the effect of walking the book, *permanent* versions, which depend in the accumulation of orders placed until time t, and *transient* versions, where the effects of past trading speeds decay with time. Good strategies are usually defined in terms of some value functional, which takes into account both the expected revenue

for the investor when employing a strategy, and some measure of the risk associated with that strategy.

5.1.2 Chapter outline

The aim of this chapter is to show how the signature method can be used to obtain approximate solutions to the problem of optimal execution. Our setting is very general, with price processes assumed to be geometric rough paths and the price impact function depends on the entire history of the trading strategy, with only a mild continuity condition assumed. The flexibility of the framework is demonstrated in part by the broad range of existing models in the literature which fall within it. An instance of this is the classical optimal execution problem presented in [26, Section 6.5], in which the underlying price process is assumed to be a Brownian motion, with an L^2 penalty imposed based on the risk of holding inventory. More recent examples include the work of Lehalle and Neumann in [93] and Cartea and Jaimungal in [24]. In [93], the authors prove results on the existence and uniqueness of an optimal trading strategy in the setting where trading signals are incorporated into the price dynamics. Similarly in [24], the authors consider the role of microstructure in the problem by including order flow as a contributing factor permanently affecting the price. Our approach can be adapted to handle models consisting of multiple correlated assets which are effected by trades in each other. Such a setting is presented in the article by Mastromatteo, Benzaquen, Eisler and Bouchaud, [107].

We begin the chapter by introducing our framework in Section 5.2. This consists of specifying our assumptions on the price process and market impact in our model, defining the space of trading speeds in which we will look for strategies, and introducing the optimal control problem. Section 5.3 is dedicated to calculating approximate solutions to the control problem. We first reformulate the problem in terms of the sig-

nature, and then we approximate the optimal trading speed by a finite-dimensional, computationally tractable minimisation problem. In Section 5.4, we provide examples of interesting extensions of the approach as it was presented in Sections 5.2 and 5.3, such as the multiple asset problem which appears in [107], and more exotic models where additional multi-dimensional noise is assumed to provide exogenous information about the price dynamics. Finally, in Section 5.5 and Section 5.6, we provide numerical evidence of the performance of the model. Good approximations to the optimal strategies in the settings [26, Section 6.5], [93] and [24] are obtained, and we also investigate the problem when the underlying price process is a fractional Brownian motion. Moreover, we demonstrate in Section 5.6 the effectiveness of our methodology on real market data.

5.2 Framework

5.2.1 Notation

Given a continuous path $X \in C([0,T];\mathbb{R})$, denote its *augmentation* by the continuous path $\widehat{X} \in C([0,T];\mathbb{R}_+ \times \mathbb{R})$ defined by $\widehat{X}_t := (t, X_t) \in \mathbb{R}_+ \times \mathbb{R}$.

Let $p \geq 1$. Define:

$$\widehat{\Omega}_T^p := \overline{\{\widehat{\mathbb{X}} \in G\Omega_p([0,T];\mathbb{R}^2) \ \mid \ X \in C^\infty([0,T];\mathbb{R}) \text{ and } X_0 = 1\}}^{d_{p-var}},$$

where the closure is taken under d_{p-var}, i.e. the p-variation distance (see [105, Definition 1.5]). Given $\widehat{\mathbb{X}} \in \widehat{\Omega}_T^p$, we will write by $X \in C([0,T];\mathbb{R})$ the unaugmented coordinate process.

Intuitively, elements of $\widehat{\Omega}_T^p$ are signatures of paths of the form (t, X_t), with initial point $X_0 = 1$. Because the first dimension of this augmented path (namely, time) is monotone increasing, and because we are only considering paths that start at 1, it follows by Corollary 2.2.7 that $\widehat{\mathbb{X}}_{0,T}^{<\infty}$ completely characterises $\widehat{\mathbb{X}}$ (and hence X).

5.2.2 The market

The space $\widehat{\Omega}_T^p$ will be our space of market paths. We will equip it with a probability space $(\widehat{\Omega}_T^p, \mathcal{B}(\widehat{\Omega}_T^p), \mathbb{P})$. Given a rough path $\widehat{\mathbb{X}} \in \widehat{\Omega}_T^p$, the unaugmented coordinate path $X : [0,T] \to \mathbb{R}$ will denote the *unaffected midprice* of the asset. In other words, X is the midprice process of the asset if the trader does not trade on the asset.

Example 5.2.1. Our market framework is very general in the sense that it includes most of the examples that have been considered in the literature. In particular, our framework includes:

1. **Continuous semimartingales**. In the literature [22, 24, 23, 93], the midprice process is often modelled as a continuous semimartingale. Continuous semimartingales can be lifted to p-geometric rough paths for $p \in (2,3)$ [104, 38, 56], and therefore they fit into our framework: the market would be given by the probability space $(\widehat{\Omega}_T^p, \mathcal{B}(\widehat{\Omega}_T^p), \mathbb{P})$ for $p \in (2,3)$ and $\mathbb{P}$ the law of the semimartingale.

2. **Fractional Brownian motion**. Our framework also includes the setting where the midprice is modelled by a fractional Brownian motion [39].

3. **Tick data**. In [52], the authors show how market tick data can be lifted to a geometric rough path, so that this example is also contained in our framework.

5.2.3 Trading speeds

In this section, we will introduce the notation of *trading speeds*.

Definition 5.2.2 (Trading speeds)**.** *Define the metrizable space* $\Lambda_T := \bigcup_{t\in[0,T]} \widehat{\Omega}_t^p$. *We define the space of trading speeds by* $\mathcal{T} := C(\Lambda_T; \mathbb{R})$. *Given a trading speed* $\theta \in \mathcal{T}$, *the trader will trade at a rate of* $\theta(\widehat{\mathbb{X}}|_{[0,t]})$.

Intuitively, the trader that is sitting at time $t \in [0,T]$ should decide how much to sell or buy by only considering what happened up to time t: she can only act based on the past, not the future. In other words, the trader's trading decision will be a (non-anticipative) function of the midprice process up to time t, i.e. $\widehat{\mathbb{X}}|_{[0,t]} \in \Lambda_T$. This intuition is incorporated into the definition of the trading speeds $\mathcal{T}$. A space similar to Λ_T was considered in [58, 36, 4, 48, 11, 120], and a similar definition of trading strategies was considered in [121].

In this chapter, the following class of trading speeds will have a special relevance:

Definition 5.2.3 (Signature trading speeds)**.** *The space of signature trading speeds $\mathcal{T}_{sig} \subset \mathcal{T}$ is defined by*

$$\mathcal{T}_{sig} := \{\theta \in \mathcal{T} \ | \ \exists \ell \in T((\mathbb{R}^2)^*) \text{ such that } \theta(\widehat{\mathbb{X}}|_{[0,t]}) = \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \ \forall \widehat{\mathbb{X}}|_{[0,t]} \in \Lambda_T\}$$

where $\widehat{\mathbb{X}}_{0,t}^{<\infty}$ denotes the signature of $\widehat{\mathbb{X}}$ over the interval $[0,t]$.

It turns out that the space of signature trading speeds $\mathcal{T}_{sig} \subset \mathcal{T}$ is very large – in fact, we have the following density result, whose proof is in Appendix B.

Lemma 5.2.4. *Let $\varepsilon > 0$. Then, there exists a compact set $\mathcal{K} \subset \widehat{\Omega}_T^p$ such that:*

1. $\mathbb{P}[\mathcal{K}] > 1 - \varepsilon$.

2. *$\mathcal{T}_{sig}$, restricted to $\mathcal{K}$, is dense in $\mathcal{T}$.*

Therefore, trading speeds can be locally approximated arbitrarily well by signature trading speeds. Hence, if one wants to optimise a certain objective function over $\mathcal{T}$, it makes sense to optimise it over $\mathcal{T}_{sig}$ instead. This is precisely the approach that will be followed in this chapter: we will look for an optimal trading speed in $\mathcal{T}_{sig}$, instead of $\mathcal{T}$.

5.2.4 Market impact

When a trader buys or sells a traded asset, the mere act of trading will affect the asset's order book. If the volume she trades is small compared to the overall volume, this effect may be neglected. However, if the trader sends large trading orders the impact on the order book may negatively affect the price at which the order is executed (see [6] and the references therein). In this section we will introduce the market impact model that will be used in this chapter.

If the trader decides to follow a signature trading speed $\theta \in \mathcal{T}_{sig}$, the *execution price* – i.e. the price the trader has access to – will be given by

$$P_t^\theta := X_t - \langle g^\theta, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle, \tag{5.1}$$

where $g^\theta \in T((\mathbb{R}^2)^*)$ is a linear functional that depends on θ that models the market impact.

Example 5.2.5. The definition of the market impact, far from being restrictive, is very general and includes many examples that have been studied in the literature. Indeed, let $\ell \in T((\mathbb{R}^2)^*)$ and set the signature trading speed $\theta(\widehat{\mathbb{X}}|_{[0,t]}) := \langle \ell, \widehat{\mathbb{X}}_{0,t} \rangle$. Then, the following are examples of market impacts included in our framework:

1. **Temporary market impact**. Set $g^\ell := \lambda \ell$, with $\lambda > 0$. Then, $\langle g^\ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle = \lambda \theta(\widehat{\mathbb{X}}|_{[0,t]})$ is the linear temporary market impact studied in [22, 24, 93]. We may also make the temporary market impact nonlinear by considering a polynomial $Q \in \mathbb{R}[x]$ and setting $g^\ell := Q^{\text{ш}}(\ell)$. Then, $\langle g^\ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle = Q(\theta(\widehat{\mathbb{X}}|_{[0,t]}))$.

2. **Permanent market impact.** In [22, 24, 23], a permanent market impact given by $\int_0^t \theta_s ds$ is considered. Setting $g^\ell := \ell \mathbf{1}$, we have $\langle g^\ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle = \langle \ell \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle = \int_0^t \langle \ell, \widehat{\mathbb{X}}|_{[0,s]} \rangle ds = \int_0^t \theta(\widehat{\mathbb{X}}|_{[0,s]}) ds$.

3. **Transient market impact**. In [61, 40, 41] the authors considered a transient market impact that is given by $\int_0^t K(t-s)\theta(\widehat{\mathbb{X}}|_{[0,s]})ds$, where $K(x) := \exp(-\rho x)$ for $\rho > 0$ constant. Then, we can find $g^\ell \in T((\mathbb{R}^2)^*)$ such that

$$\int_0^t K(t-s)\theta(\widehat{\mathbb{X}}|_{[0,s]})ds \approx \langle g^\ell, \widehat{\mathbb{X}}_{0,t}\rangle$$

to arbitrary accuracy.

4. More generally, market impacts modelled by functions of the form $G(\theta, X)$ can be well-approximated by linear functions on the signature, and they are therefore included in our framework.

5.2.5 Optimal execution problem

Suppose the trader wishes to liquidate $q_0 > 0$ units of the asset by time T. If q_0 is large compared to the traded volume, the trading activity will affect the price of the asset ([6]) negatively for the trader. Therefore, it may be more beneficial to spread the trading activity over the interval $[0, T]$ to avoid the undesired market impact. In this case, however, the trader will be exposed to market fluctuations that may affect her adversely. Hence, the task is to find a suitable trading speed to liquidate the inventory q_0 which accounts for this trade-off. We will now introduce the optimal execution problem that will be studied in this chapter.

Definition 5.2.6. *The wealth corresponding to the trading speed $\theta \in \mathcal{T}$ is defined by*

$$W_t^\theta := \int_0^t P_s^\theta \theta(\widehat{\mathbb{X}}|_{[0,s]})ds.$$

On the other hand, the remaining inventory is defined by

$$Q_t^\theta := q_0 - \int_0^t \theta(\widehat{\mathbb{X}}|_{[0,s]})ds$$

where $q_0 > 0$ is the initial inventory. We define the value function $\mathcal{V}^\theta : \widehat{\Omega}_T^p \to \mathbb{R}$ by

$$\mathcal{V}^\theta(\widehat{\mathbb{X}}) := W_T^\theta - \phi \int_0^T (Q_t^\theta)^2 dt + Q_T^\theta (P_T^\theta - \alpha Q_T^\theta) \tag{5.2}$$

with $\alpha, \phi \geq 0$ constants.

In this chapter we will study the following optimal execution problem given by the optimisation problem

$$\sup_{\theta \in \mathcal{T}} \mathbb{E}[\mathcal{V}^\theta(\widehat{\mathbb{X}})]. \tag{5.3}$$

The first term of the value function indicates that, in principle, the trader would like to maximise the wealth acquired by following the trading strategy θ. If the investor arrives the terminal time with a non-zero inventory Q_T^θ, the third term of the value function $Q_T^\theta(P_T^\theta - \alpha Q_T^\theta)$ ensures that these leftovers are executed with a penalisation $\alpha > 0$. Finally, the term $-\phi \int_0^T (Q_t^\theta)^2 dt$ penalises holding inventory for a long time. There are different interpretations for this term. For instance, this running inventory penalty could be seen as an *urgency term*. Another interpretation comes from the setting where the investor would like to account for model uncertainty: the larger ϕ is, the less certain the trader is about the dynamics imposed on the midprice (see [24, 21]). In any case, a large ϕ would increase the trading speed near the beginning, and reduce it near the end.

This particular value function was chosen due to its popularity in the literature [22, 93, 24, 61, 40, 23, 41], but the authors would like to emphasise that the methodology proposed in this chapter also applies to other alternative definitions of the value function, and we are not restricted to this particular choice of $\mathcal{V}^\theta$.

Properties of signatures, and the shuffle product property in Definition 1.3.3 in particular, will make finding the optimal trading speed for the optimal control problem (5.3) in the restricted space $\mathcal{T}_{sig} \subset \mathcal{T}$ easier to solve. Due to the density result stated

in Lemma 5.2.4, we will restrict the space of trading speeds from $\mathcal{T}$ to $\mathcal{T}_{sig}$, so that we will solve the following problem instead:

$$\sup_{\theta \in \mathcal{T}_{sig}} \mathbb{E}[\mathcal{V}^{\theta}(\widehat{\mathbb{X}})]. \tag{5.4}$$

5.3 Optimal execution

The value function (5.2) is a nonlinear function of the underlying price path. However, for signature trading strategies $\theta \in \mathcal{T}_{sig}$ it turns out to be a linear function on the signature of the midprice process. This is due to the shuffle product property (Definition 1.3.3) – each term in the value function can be rewritten as a linear function on the signature of the midprice process.

Lemma 5.3.1. *Let $\theta \in \mathcal{T}_{sig}$ be the signature trading speed given $\theta(\widehat{\mathbb{X}}|_{0,t}) = \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$, with $\ell \in T((\mathbb{R}^2)^*)$. Then, given any $\widehat{\mathbb{X}} \in \widehat{\Omega}_T^p$ and $t \in [0,T]$, we have*

1. $W_t^{\ell} = \left\langle \left((\mathbf{2} + \varnothing - g^{\ell}) \text{ш} \ell \right) \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle$.
2. $Q_t^{\ell} = \langle q_0 \varnothing - \ell \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$.
3. $\int_0^t (Q_s^{\ell})^2 ds = \langle (q_0 \varnothing - \ell \mathbf{1})^{\text{ш}2} \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$.
4. $Q_t^{\ell}(P_t^{\ell} - \alpha Q_t^{\ell}) = \langle (q_0 \varnothing - \ell \mathbf{1}) \text{ш} (\mathbf{2} + \varnothing - g^{\ell}) - \alpha (q_0 \varnothing - \ell \mathbf{1})^{\text{ш}2}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$.

Proof. Let $\widehat{\mathbb{X}} \in \widehat{\Omega}_T^p$ and $t \in [0,T]$.

1. Notice that, because $X_0 = 1$, we have $X_s = \langle \mathbf{2} + \varnothing, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle$ for each $s \in [0,t]$. Then, by the shuffle product property (Definition 1.3.3),

$$\begin{aligned} W_t^{\ell} &= \int_0^t P_s^{\ell} \langle \ell, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle ds = \int_0^t (X_s - \langle g^{\ell}, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle) \langle \ell, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle ds \\ &= \int_0^t \left\langle (\mathbf{2} + \varnothing - g^{\ell}) \text{ш} \ell, \widehat{\mathbb{X}}_{0,s}^{<\infty} \right\rangle ds = \left\langle \left((\mathbf{2} + \varnothing - g^{\ell}) \text{ш} \ell \right) \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle. \end{aligned}$$

2. Follows from the fact that $\int_0^t \langle \ell, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle ds = \langle \ell \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$.

3. Using 2.,

$$\int_0^t (Q_s^\ell)^2 ds = \int_0^t \langle (q_0 \varnothing - \ell \mathbf{1})^{\text{ш}2}, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle ds = \langle (q_0 \varnothing - \ell \mathbf{1})^{\text{ш}2} \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle.$$

4. Using 2. again,

$$\begin{aligned} Q_t^\ell (P_t^\ell - \alpha Q_t^\ell) &= \langle q_0 \varnothing - \ell \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \langle \mathbf{2} + \varnothing - g^\ell - \alpha (q_0 \varnothing - \ell \mathbf{1}), \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \\ &= \langle (q_0 \varnothing - \ell \mathbf{1}) \text{ш} (\mathbf{2} + \varnothing - g^\ell) - \alpha (q_0 \varnothing - \ell \mathbf{1})^{\text{ш}2}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle. \end{aligned}$$

□

Therefore, the optimal liquidation problem (5.3) is then transformed into the following problem:

Proposition 5.3.2. *Let $\theta \in \mathcal{T}_{sig}$ be the signature trading speed given $\theta(\widehat{\mathbb{X}}|_{0,t}) = \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$, with $\ell \in T((\mathbb{R}^2)^*)$. Then, given any $\widehat{\mathbb{X}} \in \widehat{\Omega}_T^p$ and $t \in [0,T]$, the value function can be written as*

$$\begin{aligned} \mathcal{V}^\theta(\widehat{\mathbb{X}}) = \Big\langle & \left((\mathbf{2} + \varnothing - g^\ell) \text{ш} \ell \right) \mathbf{1} - (q_0 \varnothing - \ell \mathbf{1})^{\text{ш}2} (\phi \mathbf{1} + \alpha \varnothing) + \\ & (q_0 \varnothing - \ell \mathbf{1}) \text{ш} (\mathbf{2} + \varnothing - g^\ell), \widehat{\mathbb{X}}_{0,T}^{<\infty} \Big\rangle. \end{aligned}$$

Therefore, the optimal liquidation problem (5.3) *is reduced to*

$$\begin{aligned} \sup_{\ell \in T((\mathbb{R}^2)^*)} \Big\langle & \left((\mathbf{2} + \varnothing - g^\ell) \text{ш} \ell \right) \mathbf{1} - (q_0 \varnothing - \ell \mathbf{1})^{\text{ш}2} (\phi \mathbf{1} + \alpha \varnothing) \\ & + (q_0 \varnothing - \ell \mathbf{1}) \text{ш} (\mathbf{2} + \varnothing - g^\ell), \mathbb{E} \left[\widehat{\mathbb{X}}_{0,T}^{<\infty} \right] \Big\rangle. \end{aligned} \tag{5.5}$$

The value function $\mathcal{V}^\theta(\widehat{\mathbb{X}})$ depends on two aspects: a stochastic component and the control θ. Moreover, this dependency is nonlinear. Proposition 5.3.2 separates this

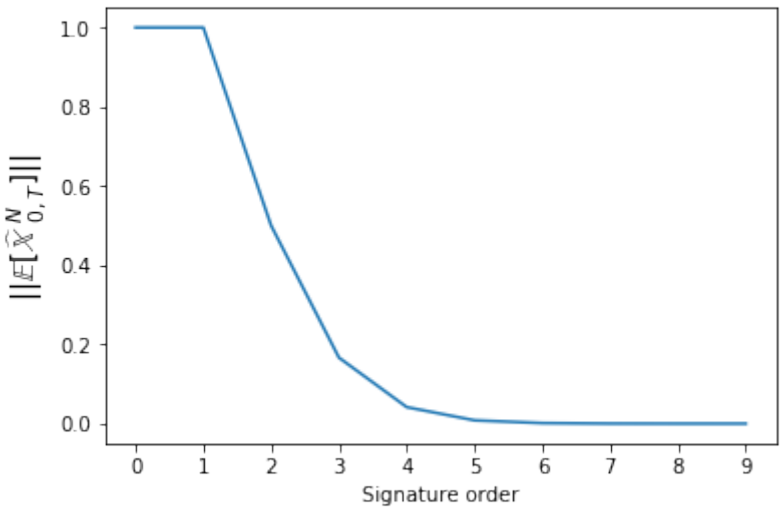

Figure 5.1: $\left\|\mathbb{E}\left[\widehat{\mathbb{X}}^N_{0,T}\right]\right\|$ as a function of N in the case where the midprice process is a Brownian motion. The factorial decay of the signature makes higher order terms small compared to the first few terms.

dependency into a deterministic component that solely depends on the control, and on a stochastic component that does not depend on the control. Moreover, because this separation makes the value function linear on the path, the expectation in (5.2) is moved inside linear functional – in other words, the resulting optimisation problem (5.5) depends on the expected signature of the midprice process.

The expected signature of the midprice process is the only dependency on the stochastic process. This object plays the analogous role of the moments of a random variable, but on path space. It was shown in fact in [33] that under certain growth assumptions, the expected signature determines the law of the stochastic process. Therefore, the fact that (5.5) depends on the expected signature of the midprice process essentially implies that the optimisation problem depends on the entire law of the process.

5.3.1 Numerically solving the optimal execution problem

The optimisation problem (5.5) from Proposition 5.3.2 involves the full expected signature $\mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{<\infty}\right]$. In practice, however, one has to consider the truncated expected signature of order $N \in \mathbb{N}$, i.e. $\mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{\leq N}\right]$.

However, the fast decay of the signature – it decays factorially – implies that the first few terms will dominate the rest, and not much information will be lost in the truncation. As a consequence, the expected signature typically also decays factorially ([33] for instance showed this fact for wide classes of Lévy, Markov and Gaussian processes).

Figure 5.1 shows $\left\|\mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{N}\right]\right\|$ plotted against N in the case where the midprice process X is a Brownian motion. As we see, the factorial decay makes higher order terms small compared to the first few terms. Therefore, in practice one doesn't need to consider truncations of very high order.

Once the signature is truncated at a certain level $N \in \mathbb{N}$, the optimisation problem (5.5) consists of finding the global maximum of a certain polynomial in several variables. For example, one can show that if a linear permanent and temporary market impact of parameters $\lambda, k > 0$ is considered, the polynomial is a concave quadratic polynomial, and finding the optimal trading speed will be reduced to finding the (unique) global maximum of this quadratic polynomial in several variables. If λ, k are negative, however, the polynomial becomes convex and its supremum is infinite – financially a negative λ and k would mean a positive market impact: the market impact benefits the trader; she will therefore benefit from trading as much as possible.

Regarding the computation of the truncated expected signature, Monte Carlo methods can be used for this task. Therefore, the only knowledge about the midprice process that is needed to solve the optimal execution problem is how to sample

from the path. The signature of a single realisation can be computed using publicly available software such as `esig`[1], `iisignature`[2] or `signatory`[3].

5.4 Extensions

In Section 5.3, we studied a certain optimal liquidation problem. In the present section we analyse different extensions of the problem, and we study how they fit in our framework.

5.4.1 Modelling the execution price with exogeneous information

For $\theta \in \mathcal{T}_{sig}$, in Section 5.2.4 the market impact was defined as a function of the trading speed and the unaffected midprocess:

$$P_t^\theta := X_t - \langle g^\theta, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle, \quad \text{with } g^\theta \in T((\mathbb{R}^2)^*). \tag{5.6}$$

However, there are other factors that affect the impact of a trading order [114, 127]. For instance, one may want to incorporate the total traded volume $V : [0,T] \to \mathbb{R}$ in the market impact [127]. Moreover, correlation and cross-asset impact between similar assets will also play a role: the execution price of an order may depend on the midprice process of other assets [114, 129, 107].

This feature can be incorporated to our framework, by modelling the execution price by

$$P_t^\theta := \langle f^\theta, \widehat{\mathbb{Z}}_{0,t}^{<\infty} \rangle, \quad \text{with } f^\theta \in T((\mathbb{R}^{n+3})^*) \tag{5.7}$$

where $\widehat{\mathbb{Z}}_{0,t}^{<\infty}$ is the signature of $\widehat{Z}_t := (t, X_t, V_t, Y_t^1, \dots, Y_t^n) \in \mathbb{R}^{n+3}$, with V_t the total traded volume up to time t and $Y_t^1, \dots, Y_t^n$ are the midprice processes of n alternative assets that the trader believes that affect the execution price of the main asset. Notice

[1] `https://pypi.org/project/esig/`
[2] [119], `https://github.com/bottler/iisignature`
[3] [88], `https://github.com/patrick-kidger/signatory`

that (5.6) is a particular case of (5.7). Other exogenous information may also be added to $\widehat{Z}$.

The methodology proposed in this chapter will then apply to this setting: the optimisation problem (5.3), for the new definition of market impact, will be reduced to an optimisation problem similar to (5.5), namely

$$\sup_{\ell \in T((\mathbb{R}^2)^*)} \left\langle \left(f^\ell \text{ш} \ell\right) \mathbf{1} - (q_0 \varnothing - \ell \mathbf{1})^{\text{ш}2} (\phi \mathbf{1} + \alpha \varnothing) + (q_0 \varnothing - \ell \mathbf{1}) \text{ ш } f^\ell, \mathbb{E}\left[\widehat{\mathbb{Z}}_{0,T}^{<\infty}\right] \right\rangle. \quad (5.8)$$

5.4.2 Optimal trading, as opposed to liquidation

In this chapter we have been focusing on the case where a trader has an initial inventory at $t = 0$, and she would like to get rid of it by time $t = T$. However, certain high-frequency traders may be interested in the following alternative question: if one starts with no inventory at $t = 0$ and one would like to finish with no inventory at $t = T$, what is the best trading strategy that can be followed on $[0, T]$? This chapter's framework can be modified for this purpose by redefining the inventory Q_t in Definition 5.2.6 by setting $q_0 = 0$.

5.4.3 Cross-asset portfolio liquidation

The discussion on Section 5.4.1 suggests another extension of the original problem studied in this chapter. Suppose there are n assets $Y^1, \ldots, Y^n$ and a trader has an initial portfolio $q = (q_1, \ldots, q_n) \in \mathbb{R}^n_+$. If the trader wishes to liquidate the inventory q (see [129, 107]), she can consider an optimal control problem similar to (5.8) that incorporates her risk profile.

More generally, the trader could aim to transition from a starting portfolio $q_{start} \in \mathbb{R}^n$ on n traded assets, to a final portfolio $q_{end} \in \mathbb{R}^n$, and she would like to do so in an optimal way. Again, our framework can be adapted for this task.

5.4.4 Other value functions

The value function considered in (5.3) was chosen in order to be consistent with the literature [22, 93, 24, 61, 40, 23, 41]. However, the methodology we propose is not intrinsic to this value function, and it can be applied to other value functions that the trader may find more appropriate.

5.5 Numerical experiments

In this section we implement the proposed methodology and test it on different settings. We first show that when we apply the methodology to various settings studied in the literature ([26, Section 6.6], [93], [24]) we retrieve the known optimal solution, thus validating the trading strategy returned by the signature methodology. Then, we apply our approach to new settings where a closed-form solution is not known.

In all the experiments in this section we begin by estimating the expected signature of the price process with Monte Carlo by sampling 50,000 path realisations of the process. Once the expected signature is estimated, we solve (5.5) to obtain a signature trading strategy – what parameters have been used in the optimisation problem will be specified on each numerical experiment. Finally, we evaluate the performance of the constructed signature trading strategy by sampling 10,000 new realisations from the price process – this will be our testing set. All performance metrics have been given on this testing set.

We have also included the time-weighted average price (TWAP) strategy [26, Section 6.3] as a baseline strategy for comparison. The TWAP strategy is a popular benchmark strategy that evenly splits the orders over the trading period.

In each experiment we provided the parameters from the optimisation problem (5.5) that were used. However, the performance of the signature method does not

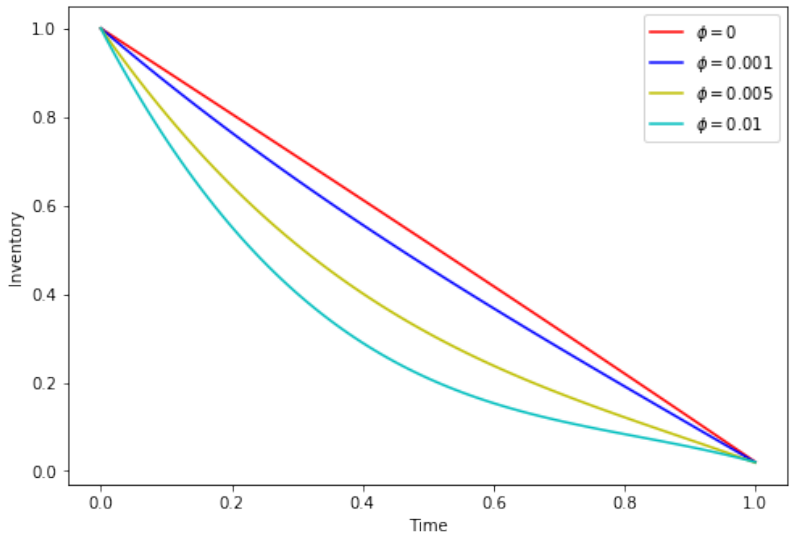

Figure 5.2: Inventory of a trader following the signature trading speed for 100 midprice path realisations and the setting considered in Section 5.5.1. Different running inventory penalties ϕ were considered. Notice that for each ϕ, the trading speed is deterministic.

appear to be sensitive with respect to the choice of these parameters in the sense that different choice of parameters also give similarly satisfactory results.

5.5.1 Brownian motion with temporary and permanent market impact

In this section we will consider the framework studied in [26, Section 6.5]. We will assume that the unaffected midprice process follows a Brownian motion with volatility σ, that is, $X_t := \sigma W_t$ with $\sigma > 0$ and W a Brownian motion. For a signature trading speed $\theta \in \mathcal{T}_{sig}$ given by $\theta(\widehat{\mathbb{X}}|_{[0,t]}) = \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$ with $\ell \in T((\mathbb{R}^2)^*)$, the execution price will be given by a *permanent* market impact and a *temporary* market impact:

$$P_t^\theta := X_t - k \int_0^t \theta(\widehat{\mathbb{X}}|_{[0,s]}) ds - \lambda \theta(\widehat{\mathbb{X}}|_{[0,t]}),$$

with $k, \lambda > 0$.

It was mentioned in Section 5.2.4 that this market impact is included in our

Theoretical optimal speed	Signature trading speed	TWAP
$0.995748 \pm 2.12 \times 10^{-5}$	$0.995697 \pm 3.97 \times 10^{-5}$	$0.993516 \pm 6.07 \times 10^{-5}$

Table 5.1: Performance of the theoretical optimal speed, signature trading speed and the time-weighted average price (TWAP) for the setting considered in Section 5.5.2. We use the expected value function, estimated using the testing set, to measure performance.

framework. More specifically, we have:

$$P_t^\ell = X_t - k\int_0^t \langle \ell, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle ds - \lambda \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle = X_t - \langle g^\ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$$

with $g^\ell := k\ell\mathbf{1} + \lambda\ell$.

We may then solve (5.5). The chosen parameters were $q_0 = 1$ share, $\lambda = 10^{-3}$ \$ $\cdot$ sec $\cdot$ share^{-1}, $k = 10^{-4}$ \$ $\cdot$ share^{-1}, $\alpha = 10$, $\sigma = 0.02$ and $T = 1$ hour, and different values for ϕ were considered. Truncated signatures of order 7 were considered to solve (5.5). As it has been established in the literature (see [26, Section 6.5]) the optimal trading speed does not depend on the midprice (i.e. it is deterministic). Moreover, if we set $\phi = 0$ so that no running inventory penalties are considered, it is known that the optimal trading speed is constant. On the other hand, when ϕ is increased, the trader decides to liquidate the inventory sooner. All this features are captured in the results we obtained – see Figure 5.2.

5.5.2 Incorporating order-flow

In [24], the authors incorporate the order-flow of all agents into the midprice dynamics. This is done by considering the midprice process

$$X_t := k\int_0^t (\mu_s^+ - \mu_s^-)ds + \sigma W_s,$$

where μ_t^+ and μ_t^- are the aggregated buying and selling orders of all market participants, respectively. These orders are assumed to follow the dynamics

$$d\mu_t^\pm = -\kappa\mu_t^\pm dt + \eta^\pm_{1+L^\pm_{t-}}\, dL_t^\pm$$

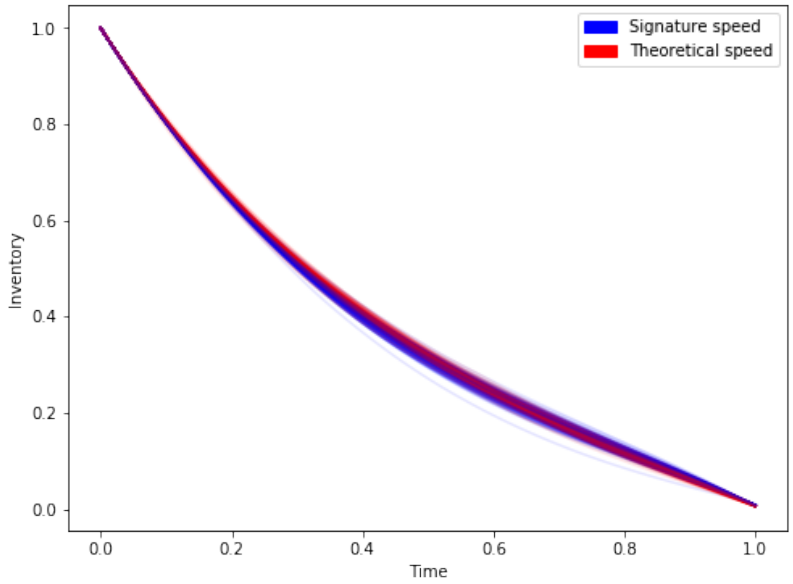

Figure 5.3: The trader's inventory for 100 midprice path realisations and the setting considered in Section 5.5.2, both for the theoretical optimal speed (red) and the signature trading speed (blue).

with $L_t^{\pm}$ independent Poisson processes of intensity λ_0, and $\eta_i^{\pm} \sim Exp(\eta_0\kappa)$ has an exponential distribution. Moreover, a temporary market impact $\lambda\theta(\widehat{\mathbb{X}}|_{[0,t]})$ was included as well.

Figure 5.3 shows the inventory for 100 realisations of the midprice path, both for the signature trading speed and the optimal trading speed that was derived in [24]. Table 5.1 shows the performance of the signature trading strategy compared to the optimal trading speed and the TWAP strategy. The parameters we considered are $\lambda = 5 \cdot 10^{-4}$ \$ $\cdot$ sec $\cdot$ share^{-1}, $k = 10^{-4}$ \$ $\cdot$ share^{-1}, $q_0 = 1$ share, $\alpha = 2$, $\phi = 5 \cdot 10^{-3}$, $\sigma = 0.1$, $\kappa = \lambda_0 = 5$, $\eta_0 = 0.8$ and signatures of order 7.

5.5.3 Incorporating trading signals

Lehalle and Neuman in [93] considered an optimal liquidation problem where the investor has access to some trading signal that predict short-term price movements, such as order book imbalance.

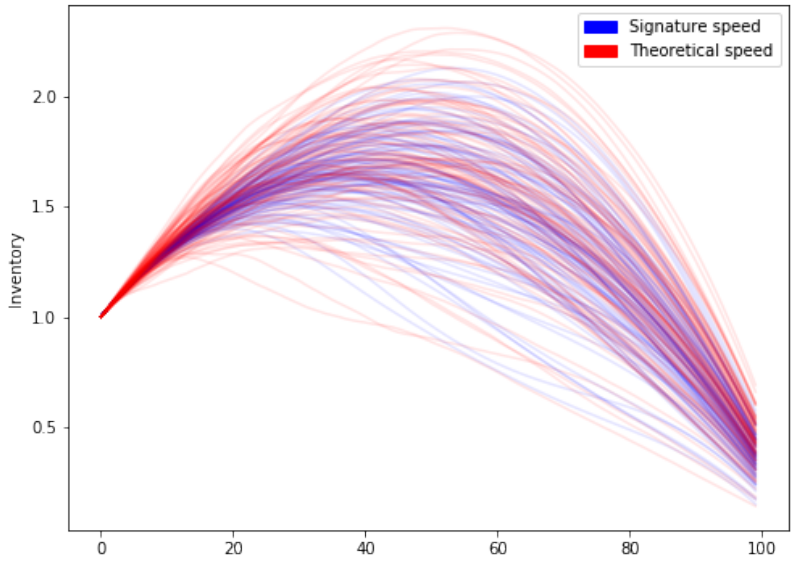

Figure 5.4: The trader's inventory for 100 midprice path realisations and the setting considered in Section 5.5.3, both for the theoretical optimal speed (red) and the signature trading speed (blue).

In this case, the midprice process was taken to be $X_t := \int_0^t I_s ds + \sigma W_t$, where I is the signal process, $\sigma > 0$ is volatility and W is a Brownian motion. In the original chapter [93], the signal I that was considered was an Ornstein-Uhlenbeck process $dI_t = -\gamma I_t dt + \sigma_0 dW_t$, where $\gamma, \sigma > 0$ are constants. Therefore, given that the midprice process is a semimartingale, this example also falls within our framework.

The price impact that was considered in [93] was a linear temporary price impact. Therefore, the execution price will be given by (D.1), where $g^\ell := \lambda \ell$ with $\lambda > 0$.

Theoretical optimal speed	**Signature trading speed**	**TWAP**
$1.0170735 \pm 4.05 \times 10^{-5}$	$1.0169903 \pm 4.12 \times 10^{-5}$	$0.999856 \pm 2.21 \times 10^{-5}$

Table 5.2: Performance of the theoretical optimal speed, signature trading speed and the time-weighted average price (TWAP) for the setting considered in Section 5.5.3. We use the expected value function, estimated using the testing set, to measure performance.

Figure 5.4 shows the running inventory for 100 realisations of the midprice process, both for the signature trading speed and the optimal trading speed that was derived in

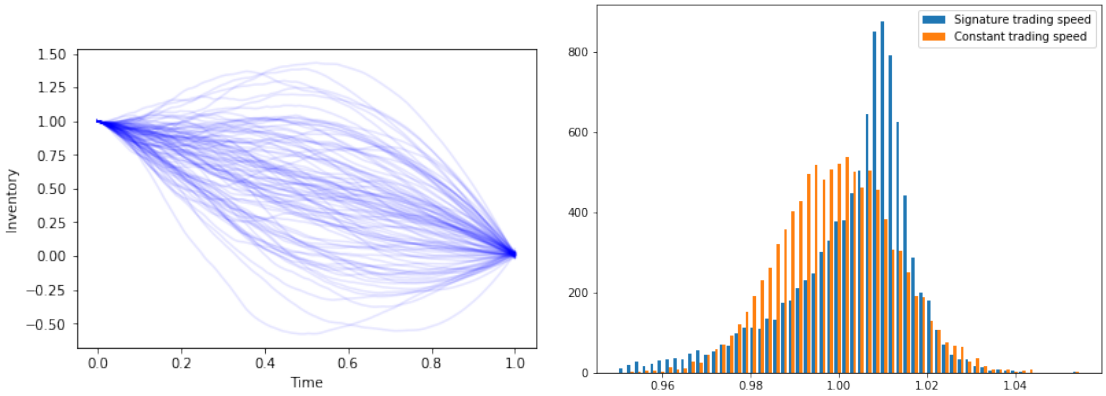

Figure 5.5: Trader's inventory (left) and the trader's wealth distribution (right) in the case where the midprice process is a fractional Brownian motion.

Signature trading speed	**TWAP**
$1.0031498 \pm 1.38 \times 10^{-4}$	$0.9991785 \pm 5.73 \times 10^{-4}$

Table 5.3: Performance of the signature trading speed and the time-weighted average price (TWAP) for the fractional Brownian motion. We use the expected value function, estimated using the testing set, to measure performance.

[93]. The chosen parameters were $q_0 = 1$ share, $\lambda = 10^{-3}$ \$·sec·share^{-1}, $\alpha = 10^{-2}, \phi = 10^{-3}, I_0 = 0.02$ and $\gamma = 0.1$. Truncated signatures of order 9 were considered. As we see, the signature trading speed seems to be a close approximation of the theoretical optimal speed. This is also reflected in the performance of the signature trading strategy with respect to the theoretical optimal speed – see Table 5.2.

Notice that the presence of the signal in the midprice process introduces a positive drift, and therefore, as illustrated by Figure 5.4, it is optimal to begin by purchasing shares in order to sell them for a profit later. This could be avoided by increasing the running inventory penalty ϕ.

5.5.4 Fractional Brownian motion

In this section, we assume that the midprice process X_t is a fractional Brownian motion. We assume a linear market impact. In other words, the execution price will

be given by

$$P_t^\theta := \sigma W_t^H - \lambda\theta(\widehat{\mathbb{X}}|_{[0,t]}),$$

where W_t^H is a fractional Brownian motion with Hurst parameter H, and $\sigma, \lambda > 0$ are constants.

Figure 5.5 shows the midprice and inventory in the case where $H = 1/3$, $\sigma = 0.02$, $q_0 = 1$ share, $\phi = 0$, $\lambda = 10^{-3}$ \$ $\cdot$ sec $\cdot$ share^{-1}, $\alpha = 0.1$, $T = 1$ hour and truncated signatures of order 7 are considered.

As we see, the behaviour differs significantly from the case where $H = 1/2$ (i.e. when X_t is a Brownian motion). Indeed, given that we don't include a running inventory penalty as $\phi = 0$, in the Brownian case we would expect the inventory Q_t to be linear. However, Figure 5.5 illustrates that this is not the case for the fractional Brownian motion, and the trading speed depends strongly on the midprice process. In fact, if we look at the expected value of the TWAP strategy (which would be optimal if the midprice process was a Brownian motion, i.e. $H = 1/2$) we see that the signature trading speed for the fractional Brownian motion outperforms TWAP (see Table 5.3). This is outperformance of the signature trading speed is reflected in the wealth distribution of both strategies shown in Figure 5.5 (right). This is not surprising, as the fractional Brownian motion with $H < 1/2$ exhibits a memory effect and an optimal trading speed should be able to exploit that.

5.6 Experiments with market data

To solve (5.5), the only information that is needed about the midprice process is its expected signature. In this section, we use real market data to estimate the expected signature, which is then used to solve (5.5). Then, we evaluated the performance of the optimal execution strategy in an out-of-sample set of market paths.

We considered midprice market data of Apple (AAPL) for 1 year, from the 1st of January 2018 to the 31st of December 2018, which was obtained from LOBSTER[4]. This data was divided into a training set of 10 months (January–October) and an out-of-sample set of 2 months (November–December).

We considered 15 minute windows from different times of each trading day – more specifically, we considered 10:00–10:15, 11:00–11:15, 12:00–12:15 and 13:00–13:15. We estimated the expected signature over each of these 15 minute windows by computing the empirical expectation of the signature (signatures of order 13 were considered) of the corresponding 15 minute windows from the testing set. Therefore, to some extent, we assume that the midprice process follows a similar behaviour over each of the windows throughout the trading year.

Once the expected signature of the midprice process for each of the 15 minute windows is estimated from the training set, we solved the optimisation problem (5.5) to estimate the optimal signature trading speed. We included a temporary and market impact:

$$P_t^\theta := X_t - k\int_0^t \theta(\widehat{\mathbb{X}}|_{[0,s]})ds - \lambda\theta(\widehat{\mathbb{X}}|_{[0,t]}).$$

The parameters we used were $\lambda = 10^{-3}$ \$ $\cdot$ sec $\cdot$ share^{-1}, $k = 10^{-4}$ \$ $\cdot$ share^{-1}, $\alpha = 0.1, \phi = 10^{-4}$ and $q_0 = 1$ share. We then evaluated the performance on the out-of-sample set for each of the 15 minute windows. Following [24], we compared the performance against the Almgren–Chriss execution strategy [3]. More specifically, we considered the *savings per share* metric (in basis points) that was used in [24], which is defined by

$$\frac{W_T - W_T^{AC}}{W_T^{AC}} \times 10^4,$$

where W_T and W_T^{AC} are the terminal wealth of the optimal signature trading speed and Almgren–Chris execution strategy, respectively.

[4] `https://lobsterdata.com/`

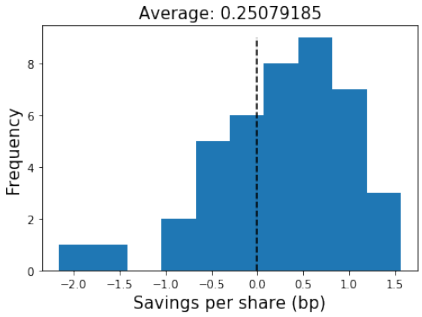

(a) 10:00–10:15.

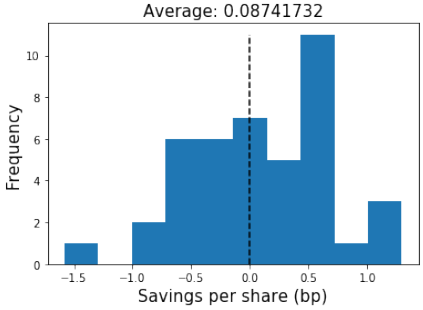

(b) 11:00–11:15.

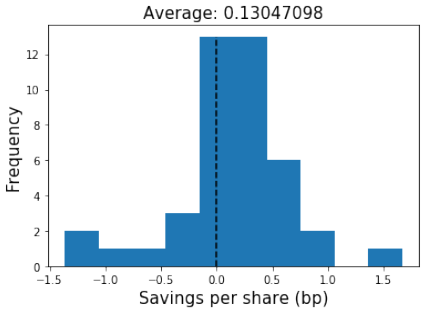

(c) 12:00–12:15.

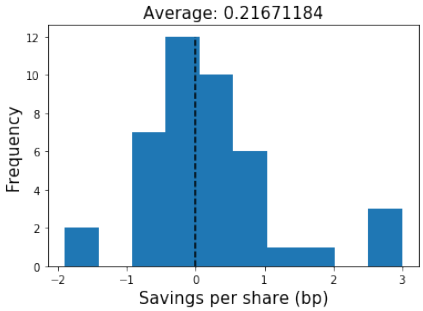

(d) 13:00–13:15.

Figure 5.6: Out-of-sample performance of the signature approach to optimal liquidation, compared to the Almgren–Chriss benchmark. The optimal signature trading speed consistently outperforms the benchmark across all 15-minute windows.

The results are shown in Figure 5.6. The optimal signature trading speed outperforms the Almgren–Chriss benchmark on all 15-minute windows, as on average the savings per share of the signature trading speed is positive.

Notice that the only assumption we have made is that the midprice process behave similarly on the same 15-minute window across different trading days. Other than that, our approach is model-free: we can, in a nonparametric and model-free way, estimate the optimal trading speed from market data.

5.7 Conclusion

In this chapter we propose a methodology to numerically approximate the solution of certain optimal execution problems. This is done in the general framework of geometric rough paths, which in particular contains many existing models in the literature.

Rough path signatures provide a methodology to reduce the original optimisation problem into a finite-dimensional, computationally feasible optimisation problem. The only information that is needed from the underlying price process is its expected signature, which can be computed using Monte Carlo methods.

This approach was tested in Section 5.5, where we show that in those cases where the optimal trading speed is known, the signature-based numerical approach is capable of retrieving it. Moreover, the generality of the approach allows the estimation of the optimal trading speed in those settings where the optimal solution is unknown. In Section 5.6 on the other hand, we showed how our methodology can be used in real market data and we demonstrated that the signature approach outperforms the Almgren–Chriss benchmark.

Chapter 6

Optimal Execution of Foreign Securities: A double-Execution Problem

The previous chapter was focused on developing a numerical method to solve certain optimal execution problems. The present chapter is more empirical in nature. Here, we aim to apply the methodology proposed in Chapter 5 to solve a certain *double-execution* problem. To the best of our knowledge, this is the first time this problem is addressed and solved in the literature. This chapter is based on the working paper [27], in collaboration with Alvaro Cartea and Leandro Sanchez Betancourt.

We employ the expected signature of equity and foreign exchange markets to derive an optimal double-execution trading strategy. The strategy maximises the wealth in the domestic currency of an investor who liquidates a block of shares in a foreign stock market and the proceeds from the sale are exchanged into the investor's domestic currency. Our approach is model agnostic because we do not specify the dynamics of the market. We prove that the optimal strategy is a linear combination of the terms in the expected signature of the market. We employ high-frequency data from Nasdaq and for various currencies to compute the signature of the market and to compare the performance of the strategies. We show that the signature-based double-execution strategy considerably outperforms various benchmarks that use single-execution strategies

(one for the shares and one for the currency). The outperformance is more pronounced for the illiquid stocks and the illiquid domestic currencies we study. The code is available in `github.com/imanolperez/optimal-double-execution`.

6.1 Introduction

In this chapter, we solve the problem of an investor who liquidates a large position in a foreign stock and the proceeds of the sale are exchanged into the investor's domestic currency. The investor only employs liquidity taking orders to trade in the limit order book (LOB) of the equity exchange and to trade in the LOB of the foreign exchange (FX). The double-execution strategy of the investor maximises the expected wealth in the domestic currency, while controlling for exposure to risks in the equity and foreign currency inventories.

We employ signature techniques to solve the double-execution problem of the investor. The signature of a path of a stochastic process (e.g., the price of a stock, the exchange rate of a currency pair, etc.) over the period $[0, t]$ is a sequence of real numbers that provides a full description of the evolution of the process up to time t. In this chapter, we show that one can write the investor's problem as a polynomial with coefficients that are linear combinations of the expected signature of the stochastic process that describes the equity and FX markets. We build on this result to train the model with high-frequency financial data and to derive the optimal double-execution strategy of the investor.

The financial data are from Nasdaq and from LMAX exchange.[1] The data are split into the 'training period' from June 2018 to December 2018 and the 'performance period' from January 2019 to June 2019. The training period is used to estimate price impact parameters and to obtain the optimal signature strategy. The performance pe-

[1] See `https://lobsterdata.com` and `https://www.lmax.com`.

riod is used to compare the performance of the double-execution strategy against two benchmarks. The benchmarks are based on the constant execution strategy known as time-weighted-average-price (TWAP). We employ data for 10 stocks traded in Nasdaq and assume that the domestic currency of the investor is one of the following: GBP, EUR, JPY, MXN, NZD. Thus, the double-execution strategy consists in selling a large position in USD and the proceeds must be exchanged into a domestic currency, all of which is done during a pre-specified trading window.

We run the investor's signature-based double-execution strategy for two hours (11:00am to 1:00pm Eastern time) every day of the performance period. In all the cases we study, the double-execution strategy considerably outperforms the benchmarks. In particular, we find that for less liquid stocks and less liquid domestic currencies, the outperformance is largest. For example, when the investor liquidates a large position in the stock of FARO Technologies (ticker FARO), the average outperformance of the double-execution strategy over the benchmarks is between 200 and 400 basis points when the domestic currency is GBP, EUR, JPY, MXN, and NZD. On the other hand, for a more liquid stock such as Apple Inc (ticker AAPL), the average outperformance is between 20 and 110 basis points for the same domestic currencies. Also, we find that the double-execution strategy performs better than the benchmark in 100% of the days for FARO, and between 85% and 96% of the days for AAPL.

Our results show that the double-execution strategy improves further when the investor computes the expected signature of the market with additional information. Specifically, the investor includes the price dynamics of a tracker of the S&P 500 index. The investor does not trade in the tracker, but the signature-based strategy employs the additional information to derive the optimal double-execution strategy. That is, the expected signature of the market includes paths from: the price of the

asset the investor trades in Nasdaq, the exchange rate of a currency pair, and the price path of the S&P 500 tracker.

We also show other applications of the double-execution strategy in currency and cryptocurrency markets. When a currency pair is illiquid, or the market does not trade a desired currency pair, the investor can employ the double-execution strategy to convert a foreign currency or a cryptocurrency into the domestic currency by trading in an intermediate currency pair. For instance, assume an investor holds a large position in Bitcoin (BTC) and must exchange it into New Zealand Dollars (NZD). However, the market does not trade the currency pair BTC/NZD, but trades the pairs BTC/USD and NZD/USD. Our findings show that a double-execution strategy that trades BTC/USD and NZD/USD to exchange BTC into NZD, considerably outperforms the TWAP-based benchmarks. Finally, we also show that an investor can avoid to trade in an illiquid currency pair by trading in other two liquid pairs that link the foreign and domestic currencies of the investor, see also [25]. For example, to exchange a large position in NZD into the Swiss franc (CHF), we show that a double-execution strategy that trades NZD/USD and USD/CHF outperforms a TWAP strategy that trades the NZD/CHF pair.

To the best of our knowledge, this is the first paper that looks at the double-execution problem faced by many investors. In addition to the applications discussed here, our framework can be widely applied and easily extended to other financial problems. Here are three additional extensions and applications of our approach. An investor: (i) employs domestic funds to go long a large position in an security (e.g., equity, bonds, derivatives) that is traded in a foreign market (e.g., buying Treasury bills using GBP); (ii) rebalances a portfolio when the proceeds from selling a long position in an asset are employed to buy into another position in the same currency (e.g., selling IBM shares to buy shares in Apple Inc.). Or, more generally, to rebalance

a portfolio of multiple assets into another portfolio of potentially different assets denominated in the same currency (e.g., rebalancing a portfolio of 50% IBM shares, 25% MSFT shares and 25% Treasury bills to a portfolio of 60% Treasury Bills, 20% IBM shares and 20% AAPL shares); (iii) includes more legs in the execution problem. For example, an investor must liquidate a position in USD to finance the purchase of equity in the Australian market. This requires to liquidate shares in the US market and exchange USD into AUD to purchase equity in the foreign market.

In the extant literature, financial problems in algorithmic trading employ classical tools of stochastic optimal control to derive trading strategies. There are two key steps in the stochastic control approach. First, to specify the law of the underlying stochastic process of the investor's problem and to write the investor's value function. Second, to show that the value function admits the dynamic programming principle and solve a Hamilton-Jacobi-Bellman (HJB) equation to obtain the trading strategies. This approach is widely used because in the cases where the trading strategies are found explicitly, one can implement them easily and draw insights that otherwise would be difficult to obtain. However, this approach has shortcomings which are relevant in our case. First, in financial markets, one hardly knows the stochastic process that underlies the investor's problem. Furthermore, in many cases, it is not straightforward to estimate or calibrate the parameters of the model. Second, often one needs to resort to numerical techniques to solve the HJB that satisfies the investor's value function – an extremely difficult task in realistic and high-dimensional cases, e.g., the double-execution problem.

Therefore, machine learning approaches are gaining popularity to solve financial problems. Machine learning is ideal in problems that: are high-dimensional, employ large data sets, require flexibility in the assumptions of the underlying stochastic process. Within machine learning, signature-based techniques are well-equipped to

solve the investor's double-execution problem because: (i) they are model agnostic – i.e., we do not specify the law of the dynamics of equity prices and of exchange rates, and (ii) because we express the investor's double-execution problem as a polynomial. This polynomial is straightforward to train with financial data and easy to employ standard optimisation packages to compute the investor's optimal signature trading strategy. Also, under some assumptions, we obtain closed-form solutions for the double-execution strategy.

In recent years, signatures have been used in various areas of mathematical finance. For instance, [102, 101] employ signatures to price and hedge exotic options and [86] to solve optimal execution problems. The work of [74] uses signatures to extract features from financial data and that of [99] to reverse-engineer trading strategies from the returns they generate.

The remainder of this chapter proceeds as follows. Section 6.2 presents the double-execution model of the investor and Section 6.3 shows how to employ the signature of the market to solve the investor's optimisation problem. Section 6.4 describes the high-frequency equity and FX data we employ and presents the estimates of the price impact parameters. Section 6.5 employs the out-of-sample data to show the performance of the double-execution strategy. Section 6.6 discusses the use of signatures and neural networks. Finally, Section 6.7 concludes and in the appendix we provide further information about the theory of signatures of paths and collect some proofs.

6.2 Model: Double-Execution

Let $\boldsymbol{P}_t = (P_t^1,\, P_t^2)$ be a pair of stochastic processes, where P_t^1 represents the best bid price of the stock and P_t^2 represents the best bid exchange rate of the currency pair

to trade units of the domestic currency in exchange for units of the foreign currency.[2] We impose the technical assumption that $\widehat{X}_t := (t, \boldsymbol{P}_t)$ admits a canonical geometric rough path lift almost surely (see Definition 2.2.2). The investor's speed of trading is denoted by $\boldsymbol{\theta}_t = (\theta_t^1, \theta_t^2)$, where θ_t^1 is the speed of trading in the stock and θ_t^2 is the speed of trading in the currency pair. When the speed of trading is positive (negative), the agent sells (buys) the security.[3] For example, assume the speed $\theta_t^1 > 0$ is fixed over a time interval of length Δt, then over that time interval the investor sells $\theta_t^1 \, \Delta t$ shares. Finally, let $(\Omega, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ be a probability space, where the filtration is the natural filtration generated by $\boldsymbol{P}$.

The trading window is $[0, T]$, where $T > 0$ is the terminal date of the execution programme in both equity and FX. The market orders sent by the investor have permanent and temporary price impact. We assume that the permanent price impact is linear in the speed of trading, so the prices and exchange rates are given by

$$P_t^{\theta^i} = P_t^i - b_i \int_0^t \theta_s^i \, \mathrm{d}s \, . \tag{6.1}$$

Here, $b_i \geq 0$ is the permanent price impact parameter, where $i = 1$ denotes equity and $i = 2$ denotes a currency pair in the FX market.

The market orders sent by the investor to the exchange receive worse prices than those of the best quotes because the market orders walk the LOB – we refer to this as temporary price impact. Specifically, when the investor sells shares and sells the

[2]Note that we make the convention that the investor 'sells' the currency pair to exchange the foreign currency into the domestic currency.

[3]By definition, 'speed' cannot be negative, but here, to be consistent with the extant literature we assume it can take on positive and negative values to denote buying and selling activity.

currency pair,[4] the execution prices the investor receives are

$$\widetilde{P}_t^{\theta^i} = P_t^{\theta^i} - k_i\,\theta_t^i\,, \tag{6.2}$$

where $k_i \geq 0$ is the temporary price impact parameter, $i \in \{1, 2\}$. We refer to this effect as temporary because we assume that the LOB replenishes immediately after the investor's trades are filled, see [26].

We denote the inventories in the foreign stock and foreign currency by $\boldsymbol{Q}_t^{\boldsymbol{\theta}} = (Q_t^{\theta^1}, Q_t^{\theta^2})$ with initial condition $\boldsymbol{Q}_0^{\boldsymbol{\theta}} = (q_0, 0)$, where $q_0 > 0$ is the amount of shares the investor must liquidate in the foreign equity market. The inventory in the foreign stock satisfies

$$Q_t^{\theta^1} = q_0 - \int_0^t \theta_s^1\,\mathrm{d}s \tag{6.3}$$

and the inventory in the foreign currency satisfies

$$Q_t^{\theta^2} = \int_0^t \theta_s^1\,\widetilde{P}_t^{\theta^1}\mathrm{d}s - \int_0^t \theta_s^2\,\mathrm{d}s\,. \tag{6.4}$$

The first term on the right-hand side of (6.4) denotes the proceeds that the investor receives in the foreign currency from selling the foreign stock. The second term on the right-hand side of (6.4) denotes the amount in the foreign currency that is exchanged into the domestic currency.

The investor's wealth in the domestic currency is

$$W_t^{\theta^2} = \int_0^t \widetilde{P}_s^{\theta^2}\,\theta_s^2\,\mathrm{d}s\,, \tag{6.5}$$

[4]A currency pair base/quote, exchanges units of the quote currency for units of the base currency. The ask side of the LOB for a currency pair shows the liquidity that market makers offer to give units of the base currency in exchange of units of the quote currency. The bid side of such pair displays the liquidity that market makers offer to receive units of the base currency in exchange for units of the quote currency. In this chapter we take the quote currency to be the final currency in which the investor wishes to have her wealth, therefore, we are interested in the bid side of the LOB for the base/quote currency pair.

and the performance criterion of the investor is

$$\mathbb{E}\big[\mathcal{V}^{\boldsymbol{\theta}}\big] \;=\; \mathbb{E}\Big[W_T^{\theta^2} + Q_T^{\theta^1}\, P_T^{\theta^1}\, \Big(P_T^{\theta^2} - \alpha_1\, Q_T^{\theta^1}\, P_T^{\theta^1}\Big) + Q_T^{\theta^2}\, \Big(P_T^{\theta^2} - \alpha_2\, Q_T^{\theta^2}\Big) \quad (6.6a)$$

$$-\phi_1 \int_0^T \Big(Q_t^{\theta^1}\, P_t^{\theta^1}\, P_t^{\theta^2}\Big)^2 \,\mathrm{d}t - \phi_2 \int_0^T \Big(Q_t^{\theta^2}\, P_t^{\theta^2}\Big)^2 \,\mathrm{d}t\Big], \quad (6.6b)$$

where $\mathbb{E}[\,\cdot\,]$ is the expectation operator and the penalty parameters $\alpha_1,\, \alpha_2\,, \phi_1,\, \phi_2$ are non-negative constants.

We comment on the various terms of the performance criterion. The first term on the right-hand side of (6.6a) is the terminal wealth in domestic currency. The second and third terms on the right-hand side of (6.6a) represent the liquidation of the final inventory positions, which includes a liquidation penalty. For example, if by the terminal date there are $Q_T^{\theta^1} > 0$ shares outstanding, then the investor liquidates the shares at the best quote in the stock market and exchanges the foreign currency into the domestic currency. These two transactions pay a penalty that includes the costs of walking the equity and the FX books. The value of the terminal penalty for the equity leg of the transaction is $\alpha_1\,(Q_T^{\theta^1}\, P_T^{\theta^1})^2$, where the units of the parameter α_1 are such that this penalty is denominated in the domestic currency. The interpretation of the third term is similar – the terminal inventory in the foreign currency is exchanged into the domestic currency, and the transaction pays a penalty that increases in the size of the inventory.[5]

The last two terms in (6.6b) represent running inventory penalties. These two penalties are not financial, i.e., the investor's cash position is not affected by these two terms. However, the strategy to liquidate the position in the stock and in the foreign currency depends on the severity of these penalties. When the value of the running inventory parameter ϕ_i is high, the speed of liquidation of the inventory in i increases. The parameters ϕ_i are also referred to as urgency parameters: a high

[5]When $\alpha_i \gg k_i$ the investor seeks optimal strategies that ensure full liquidation of the inventories in equity and the stock by the terminal date of the trading window.

value of the parameter ϕ_i indicates urgency to liquidate a large proportion of the initial inventory in i very quickly. The work of [21] shows that the running inventory penalties are equivalent to ambiguity aversion (Knightian uncertainty) to the drift of the prices and exchange rates in $\boldsymbol{P}$.

The investor's objective is to find the optimal double-execution strategy that maximises the performance criterion in (6.6), i.e., solve

$$\sup_{\boldsymbol{\theta}} \mathbb{E}\left[\mathcal{V}^{\boldsymbol{\theta}}\right] . \tag{6.7}$$

In the next section we describe how we employ machine learning and signature techniques to solve the investor's problem.

6.3 Machine learning and signatures

In recent years, with the advent of large data sets and the increase in computational power, there has been a surge in the use of machine learning techniques to solve problems that arise in finance. A popular choice is to employ neural networks (NNs) to approximate trading strategies in optimal hedging problems for financial derivatives and to derive algorithmic trading strategies; see for example [84], [13], [113], [28], [73].

The use of NNs is becoming a popular tool to solve financial problems for several reasons. First, NNs are good to approximate functions. In financial applications, NNs are employed to obtain 'approximate' strategies in which the input is the state of the market and the output is, for example, a trading or hedging strategy. In the case of execution problems, NNs are used to approximate the best strategy to buy and sell the security. Second, the literature that looks at optimisation (i.e., training) of NNs is vast– it is much simpler to optimise over a class of NNs-based trading strategies than over the broader class of all admissible trading strategies, see e.g., [62].

In this chapter, we use machine learning techniques to solve the double-execution problem in (6.7). We approximate trading strategies with signature-based strategies,

see [86]. As with NNs, signature-based trading strategies can approximate trading strategies and it is straightforward to train the signature-based strategies to optimise the performance criterion of the investor.

One of the key steps to solve the double-execution problem in (6.7) is to compute the expected signature of the price process $\boldsymbol{P}$. We also refer to the process $\boldsymbol{P}$ as the market because all the market information the investor employs is encoded in $\boldsymbol{P}$ – below in Section 6.6 we discuss further the implementation of NNs and signatures. Below we discuss how to augment the dimension of $\boldsymbol{P}$ to increase the information of the signature of the market and in Section 6.5 we show how the performance of the double-execution strategy improves with additional information of a third traded asset.

We do not follow the approach in the classical algorithmic trading literature in which one assumes to know the law of the stochastic process $\boldsymbol{P}$ and its parameters are estimated with market data. Our approach is model agnostic. In the following subsection we show how to express the double-execution problem of the investor in terms of a signature-based trading strategy and as a function of a finite number of terms of the expected signature of the market.

6.3.1 Signature trading strategies

Recall that the investor wishes to liquidate a large position in a foreign asset and exchange the proceeds of the sales into her domestic currency. All the trades are aggressive market orders, so the investor is interested in the bid prices in the security and the bid or ask exchange rates in the FX market; for simplicity we assume that the investor's orders in the FX are filled by the bid side of the FX LOB. At time $t \in [0, T]$, the investor has historical information about the path of the best bid prices and best bid exchange rates from 0 to t. This historical information is encoded in the

signature of $\boldsymbol{P}$, which is a sequence of statistics that describes the historical evolution of prices over the period $[0,t]$, e.g., the trend, realised volatility, etc., see [100]. It is well-known that the evolution of $\boldsymbol{P}$ over $[0,t]$ is uniquely determined by its realised signature, see [15].

For each path of $\boldsymbol{P}$, denote $\widehat{X}_t := (t, \boldsymbol{P}_t) \in \mathbb{R} \oplus \mathbb{R}^2$ as in Chapter 5. We then denote the signature of $\widehat{X}$ (recall Definition 2.2.4) by $\widehat{\mathbb{X}}_{0,t}^{<\infty}$ for each $t \in [0,T]$.

In general, investors seek trading strategies that optimise a performance criterion in a stochastic environment. When trading strategies are available in closed-form, it is clear to see how the strategies depend on the current state of the system and, in most cases, also depend on future realisations of the stochastic process that underlies the financial problem.[6]

In our approach, instead of numerically solving a non-linear PDE to obtain the investor's optimal policy (which also requires a specific model for the dynamics of $\boldsymbol{P}$), or instead of employing NNs to approximate the trading strategies, we follow Chapter 5 and use the information encoded in the signature of the stochastic process that underlies the optimisation problem to obtain the optimal trading strategies. This is done via the signature trading strategies used in Chapter 5. In our case, all relevant information that drives the uncertainty in the investor's problem is given by the signature of the stochastic process $\boldsymbol{P}$.

Definition 6.3.1 (Class of signature trading strategies)**.** *A double-execution speed* $\boldsymbol{\theta}_t = (\theta_t^1, \theta_t^2)$ *is a signature trading strategy if for all* $t \in [0,T]$

$$\theta_t^1 = \langle \beta, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \,, \tag{6.8a}$$

$$\theta_t^2 = \langle \gamma, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \,, \tag{6.8b}$$

[6]There is a large body of literature on algorithmic trading. See for example: [72, 10, 94, 69, 24, 71].

where $\beta \in T^{(N)}((\mathbb{R}^3)^*)$, $\gamma \in T^{(M)}((\mathbb{R}^3)^*)$ *are linear functionals (the strategy's parameters) and* $N, M \in \mathbb{N}$ *are hyperparameters.*

The class of signature trading strategies is large. For example, they can approximate arbitrarily well general trading strategies (recall Lemma 5.2.4). Therefore, we optimise the double-execution problem in (6.7) over the class of signature trading strategies in (6.8), instead of the larger class of general trading speeds.

6.3.2 Expected signature of the market

Ideally, the investor knows the law of the process $\boldsymbol{P}$, however this is hardly the case in problems that involve financial data and the dynamics of financial markets. In our setup, the investor only knows the past evolution of the stochastic process $\boldsymbol{P}$ and employs historical observations to estimate the expected signature of $\boldsymbol{P}$. The investor can augment the market information employed to solve the double-execution problem by increasing the dimension of $\boldsymbol{P}$ with other features such as: traded volumes, volume imbalance in the LOB, prices of other assets, order flow, etc.

The expected signature of the market is the infinite sequence

$$\mathbf{E} := \mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{<\infty}\right]. \tag{6.9}$$

For each word $\mathbf{w} \in \mathcal{W}(\mathcal{A}_3)$ (Section 1.3) we denote

$$\mathbf{E}^{\mathbf{w}} := \langle \mathbf{w}, \mathbf{E} \rangle = \mathbb{E}\left[\langle \mathbf{w}, \widehat{\mathbb{X}}_{0,T}^{<\infty} \rangle\right].$$

6.3.3 Double-execution problem

Here, we express the agent's expected terminal wealth and the double-execution problem in (6.7) for a class of signature trading strategies. We fix the value of $M, N \in \mathbb{N}$ and let $\boldsymbol{\theta}$ denote a signature trading strategy with coefficients $\beta \in T^{(N)}((\mathbb{R}^3)^*)$ and

$\gamma \in T^{(M)}((\mathbb{R}^3)^*)$ as in Definition 6.3.1, and write β and γ in terms of the basis of words $\mathcal{W}(\mathcal{A}_3)$:

$$\beta = \sum_{\substack{\mathbf{w}\in\mathcal{W}(\mathcal{A}_3)\\ |\mathbf{w}|\leq N}} \beta_{\mathbf{w}}\mathbf{W}, \qquad \gamma = \sum_{\substack{\mathbf{w}\in\mathcal{W}(\mathcal{A}_3)\\ |\mathbf{w}|\leq M}} \gamma_{\mathbf{w}}\mathbf{W}$$

with $\beta_{\mathbf{w}}, \gamma_{\mathbf{w}} \in \mathbb{R}$.

After tedious, but straightforward calculations, one can show that the expected terminal wealth is a quadratic polynomial in the parameters $\gamma_{\mathbf{w}}$, whose coefficients are linear combinations of a finite number of terms of the expected signature of the market:

$$\mathbb{E}[W_T] = \sum_{\substack{\mathbf{w}\in\mathcal{W}(\mathcal{A}_3)\\ |\mathbf{w}|\leq M}} \left(\sum_{\mathbf{z}\in\mathcal{W}(\mathcal{A}_3)} c_{\mathbf{z}}^{\mathbf{w}}\, \mathbf{E}^{\mathbf{z}} \right) \gamma_{\mathbf{w}} + \sum_{\substack{\mathbf{v},\mathbf{w}\in\mathcal{W}(\mathcal{A}_3)\\ |\mathbf{w}|\leq M}}^{M} \left(\sum_{\mathbf{z}\in\mathcal{W}(\mathcal{A}_3)} c_{\mathbf{z}}^{\mathbf{w},\mathbf{v}}\, \mathbf{E}^{\mathbf{z}} \right) \gamma_{\mathbf{w}}\, \gamma_{\mathbf{v}}\,, \quad (6.10)$$

where $c_{\mathbf{z}}^{\mathbf{w}}, c_{\mathbf{z}}^{\mathbf{w},\mathbf{v}} \in \mathbb{R}$ are known coefficients that do not depend on the trading strategy $\boldsymbol{\theta}$ or on the market's expected signature $\mathbf{E}$, and we note that most coefficients $c_{\mathbf{z}}^{\mathbf{w},\mathbf{v}}$ are zero.

Similarly, for the same class of signature strategies $\boldsymbol{\theta}$, the theorem below shows that the double-execution problem in (6.7) is a polynomial of degree six in the variables $\beta_{\mathbf{w}}, \gamma_{\mathbf{v}}$. The coefficients of this polynomial are linear combinations of a finite number of terms of the expected signature of the market.

Theorem 6.3.2. *Let $\boldsymbol{\theta}$ be a signature trading strategy with coefficients*

$$\beta = \sum_{\substack{\mathbf{w}\in\mathcal{W}(\mathcal{A}_3)\\ |\mathbf{w}|\leq N}} \beta_{\mathbf{w}}\mathbf{W}, \qquad \gamma = \sum_{\substack{\mathbf{w}\in\mathcal{W}(\mathcal{A}_3)\\ |\mathbf{w}|\leq M}} \gamma_{\mathbf{w}}\mathbf{W}.$$

Then, there exists an explicit shuffle polynomial (recall Definition 1.3.5) $X^{ш}(\boldsymbol{\beta}, \boldsymbol{\gamma})$ of order six such that the performance criterion in (6.6) *is given by*

$$\mathbb{E}\left[\mathcal{V}^{\boldsymbol{\theta}}\right] = \langle X^{ш}(\beta, \gamma), \mathbf{E}\rangle\ .$$

A closed-form expression of $X^{\text{ш}}$ is given in Proposition D.0.2 in the appendix.

Equivalently, there are coefficients $c_{n,k}^{i_1,\dots,i_{d-k},j_1,\dots,j_k} \in \mathbb{R}$ (which do not depend on $\boldsymbol{\theta}$ or $\mathbf{E}$) that depend on the penalty parameters $\alpha_1, \alpha_2, \phi_1, \phi_2$, such that the performance criterion is written as

$$\mathbb{E}\left[\mathcal{V}^{\boldsymbol{\theta}}\right] = \sum_{d=1}^{6}\sum_{k=0}^{d} \sum_{\substack{\mathbf{w}_1,\dots,\mathbf{w}_{d-k}\in\mathcal{W}(\mathcal{A}_3)\\ |\mathbf{w}_1|,\dots,|\mathbf{w}_{d-k}|\leq N}} \sum_{\substack{\mathbf{w}_1,\dots,\mathbf{v}_k\in\mathcal{W}(\mathcal{A}_3)\\ |\mathbf{v}_1|,\dots,|\mathbf{v}_{d-k}|\leq N}} \left(\sum_{\mathbf{z}\in\mathcal{W}(\mathcal{A}_3)} c_{\mathbf{z},k}^{\mathbf{w}_1,\dots,\mathbf{w}_{d-k},\mathbf{v},\dots,\mathbf{v}_k} \mathbf{E}^{\mathbf{z}} \right) \beta_{\mathbf{w}_1}\cdots\beta_{\mathbf{w}_{d-k}}\gamma_{\mathbf{v}_1}\cdots\gamma_{\mathbf{v}_k}. \tag{6.11}$$

Therefore, the double-execution problem in (6.7) *reduces to the following optimisation problem:*

$$\max_{\substack{\beta\in T^{(N)}((\mathbb{R}^3)^*),\\ \gamma\in T^{(M)}((\mathbb{R}^3)^*)}} \sum_{d=1}^{6}\sum_{k=0}^{d} \sum_{\substack{\mathbf{w}_1,\dots,\mathbf{w}_{d-k}\in\mathcal{W}(\mathcal{A}_3)\\ |\mathbf{w}_1|,\dots,|\mathbf{w}_{d-k}|\leq N}} \sum_{\substack{\mathbf{w}_1,\dots,\mathbf{v}_k\in\mathcal{W}(\mathcal{A}_3)\\ |\mathbf{v}_1|,\dots,|\mathbf{v}_{d-k}|\leq N}} \left(\sum_{\mathbf{z}\in\mathcal{W}(\mathcal{A}_3)} c_{\mathbf{z},k}^{\mathbf{w}_1,\dots,\mathbf{w}_{d-k},\mathbf{v},\dots,\mathbf{v}_k} \mathbf{E}^{\mathbf{z}} \right) \beta_{\mathbf{w}_1}\cdots\beta_{\mathbf{w}_{d-k}}\gamma_{\mathbf{v}_1}\cdots\gamma_{\mathbf{v}_k}. \tag{6.12}$$

For a proof see Proposition D.0.2 in the appendix.

The number of terms of the expected signature of the market that appear in the optimisation problem in (6.12) depends on the value of the hyperparameters M and N; the higher are the values of M and N, the more terms of the expected signature appear in the performance criterion (6.11). When the value of M, N are finite, the expected signature of the market (6.12) contains a finite number of terms.

In applications, the choice of the value of the hyperparameters M and N is similar to the choice of the number of hidden layers or neurons in a NN. When the values of M and N are large, the class of approximate trading strategies in (6.8) is more exhaustive. However, this would also increase the dimensionality of the polynomial optimisation in (6.12), which makes it more difficult to solve; we return to this point below in Section 6.6.

In the extreme case where $M = N = 0$, the signature trading strategy in Definition 6.3.1 is constant because θ^1 and θ^2 are determined by the coefficients $\beta_\varnothing, \gamma_\varnothing \in \mathbb{R}$ and

$\langle \varnothing, \widehat{\mathbb{X}}^{<\infty}_{0,t} \rangle = 1$, i.e., the strategy does not depend on the realised signature of the market during the trading window.

Proposition 6.3.3. *Let $\phi_1 = \phi_2 = b_1 = b_2 = 0$. Then, the polynomial in* (6.12) *reduces to a polynomial of degree four in the variables β and γ, and the double-execution problem becomes*

$$\begin{aligned}\max_{\beta,\gamma\in\mathbb{R}} \quad & \Big(\mathbf{E^{21}} + \mathbf{E^{213}} + \mathbf{E^{231}} + \mathbf{E^{321}} - (\mathbf{E^{2}} + \mathbf{E^{23}} + \mathbf{E^{32}})\, T + 2\,(1 + 2\,\left(\mathbf{E^{2}} + \mathbf{E^{22}}\right))\; q_0\, T\, \alpha_1\Big)\,\beta \\ & - \Big((1 + \mathbf{E^{3}})\, k_1\, T + (1 + 2\,\mathbf{E^{2}} + 2\,\mathbf{E^{22}})\, \alpha_1\, T^2 + \left(2\,\mathbf{E^{2121}} + 4\mathbf{E^{2211}} + 2\,\mathbf{E^{21}}\, T + T^2\right)\; \alpha_2\Big)\,\beta^2 \\ & + \left(\mathbf{E^{21}} + T\right)\; 2\,\alpha_2\, k_1\, T\,\beta^3 - k_1^2\, T^2\, \alpha_2\,\beta^4 + \left(\mathbf{E^{31}} - \mathbf{E^{3}}\, T\right)\; \gamma - \left(k_2\, T + T^2 \alpha_2\right)\; \gamma^2 \\ & + \left(\mathbf{E^{21}} + T\right)\; \alpha_2\, 2\, T\,\beta\,\gamma - 2\, k_1\, T^2\, \alpha_2\,\gamma\,\beta^2\,. \end{aligned} \tag{6.13}$$

Proof. Use expression (D.3) in the appendix with $M = N = 0$. □

We remark some of the features of the constant double-execution strategy in Proposition 6.3.3. First, we see that the coefficients of the polynomial that have both β_0 and γ_0 are all linear in the terminal penalty parameter α_2. Thus, for strategies that do not depend on the realised signature of the market, the cross dependence of the speed of trading in the equity and FX markets is via the penalty to liquidate terminal inventory in the FX market. Note, however, that the terms of the expected signature of the market that appear in the polynomial above depend on the dynamics of the price of the stock and the dynamics of the exchange rate.

The following proposition shows that when the terminal liquidation penalty parameter in the FX market is zero (i.e., $\alpha_2 = 0$), the constant strategies are straightforward to derive and we obtain, under further assumptions, TWAP strategies in equity, and strategies that take speculative positions in the equity and the FX markets.

Corollary 6.3.4. *Assume the double-execution of the investor is that in Proposition 6.3.3 and let $\alpha_2 = 0$. Then, for $t \in [0, T]$, the optimal speeds of trading in the equity*

and FX markets are

$$\theta_t^{1,*} = \frac{\mathbf{E^{21}} - (\mathbf{E^{2}} + \mathbf{E^{23}} + \mathbf{E^{32}})\, T + \mathbf{E^{231}} + \mathbf{E^{321}} + \mathbf{E^{213}} + 2\,(1 + 2\,\mathbf{E^{2}} + 2\,\mathbf{E^{22}})\, q_0\, T\, \alpha_1}{2\,(k_1\, T + \mathbf{E^{3}}\, k_1\, T + (1 + 2\,\mathbf{E^{2}} + 2\,\mathbf{E^{22}})\, T^2\, \alpha_1)} \tag{6.14}$$

and

$$\theta_t^{2,*} = \frac{\mathbf{E^{31}} - \mathbf{E^{3}}\, T}{2\, k_2\, T}\,. \tag{6.15}$$

Replacing each term of the expected signature with its respective value, the optimal speeds $\theta_t^{1,\star}$ *and* $\theta_t^{2,\star}$ *in* (6.14) *and* (6.15) *are simplified to*

$$\theta_t^{1,*} = \frac{2\,\alpha_1\, q_0\, \mathbb{E}\left[(P_T^1)^2\right] - \mathbb{E}\left[P_T^2\, \left(P_T^1 - \frac{1}{T}\int_0^T P_s^1\, \mathrm{d}s\right)\right]}{2\,\alpha_1\, T\, \mathbb{E}\left[(P_T^1)^2\right] + 2\, k_1\, \mathbb{E}\left[P_T^2\right]} \tag{6.16}$$

and

$$\theta_t^{2,*} = -\frac{\mathbb{E}\left[P_T^2\right] - \frac{1}{T}\int_0^T \mathbb{E}\left[P_s^2\right]\, \mathrm{d}s}{2\, k_2}\,. \tag{6.17}$$

Proof. Compute first order conditions with respect to β, γ and note that the signature strategy is given by $\theta^{1,*} = \beta^*$ and $\theta^{2,*} = \gamma^*$. The second order conditions are trivial to check. □

We comment on the results of the corollary above. We look at various cases where the speed reduces to simple strategies, such as TWAP, and discuss how the strategy speculates on expected moves in prices. First, we focus on the speed of trading in the equity market for two extreme cases of the value of the terminal penalty parameter α_1. Then, we look at how the strategy in equity takes speculative positions that depend on the expected value of terminal inventory in the domestic currency, after which we look at the effect of volatility of the price of equity on the optimal strategy. Finally, we discuss the constant speed of trading in the FX market, which is based solely on speculative trades.

Case 1: P^1, P^2 are independent and $\alpha_1 = 0$. The speed in (6.16) becomes

$$\theta_t^{1,*} = -\frac{\mathbb{E}\left[P_T^1\right] - \frac{1}{T}\int_0^T \mathbb{E}\left[P_s^1\right]\, \mathrm{d}s}{2\, k_1}\,, \tag{6.18}$$

which is a constant strategy that speculates on 'the trend' of the price of equity. Here, the trend refers to the quantity $\mathbb{E}\left[P_T^1\right] - \frac{1}{T}\int_0^T \mathbb{E}\left[P_s^1\right] \mathrm{d}s$. For example, if the expected value of the price of equity at time T is higher (lower) than the average of the expected price of the equity over the interval $[0, T]$, the trend is positive (negative) and the optimal strategy is to buy (sell) equity at the speed in (6.18) during the entire trading horizon – recall that when $\theta_t > 0$ ($\theta_t < 0$) the investor sells (buys). In other words, when the trend is positive, the equity will be trading (on average) at a lower price than the expected price at time T; thus the investor speculates: buys equity at low prices and unwinds the position at a higher price. Clearly, the higher the costs of walking the equity book (i.e., the higher is the value of the impact parameter k_1), the smaller are the speculative positions the investor takes. Note that the speculative position in equity will be closed at time T without paying a penalty and without walking the book because we assume $\alpha_1 = 0$.

Case 2: Terminal penalty parameter $\alpha_1 \to \infty$. The speed in (6.16) becomes

$$\theta^{1,*} = \frac{q_0}{T}\,, \tag{6.19}$$

which is TWAP. The intuition is as follows. First, note that if the value of the terminal inventory in equity q_T is different from zero, the investor would pay an infinite penalty in the domestic currency. Thus, the investor needs to search over all constant strategies (because we assumed that $M = N = 0$) that guarantee full liquidation by the end of the trading horizon, while minimising the costs of walking the equity LOB. This strategy is TWAP.

Case 3: Speculation in equity. We focus on the term

$$\mathbb{E}\left[P_T^2\left(P_T^1 - \frac{1}{T}\int_0^T P_s^1 \,\mathrm{d}s\right)\right] \tag{6.20}$$

that appears in (6.16). This term is the average expected profit from speculating at a constant speed during the trading horizon. Everything else being equal, its

contribution to the double-execution strategy is to send speculative trades to buy or sell equity (if the term is zero there are no speculative trades). The intuition is as follows. For simplicity, let $q_0 = 0$. Assume the investor trades at the constant speed ν during $[0, T]$, and at time T the investor closes the position. The profit in the foreign currency is $\nu\, T\, P_T^1 - \nu \int_0^T P_s^1\, ds$ and in the domestic currency the profit is

$$P_T^2 \left(\nu\, T\, P_T^1 - \nu \int_0^T P_s^1\, ds \right) . \tag{6.21}$$

Next, divide the profits in (6.21) by the total number of shares traded $\nu\, T$ and take expectations to obtain the expression in (6.20). This expression is the average of the time zero expected profit from speculating at a constant speed during the trading horizon $[0, T]$. Therefore, the investor buys (sells) equity when the expected value of the speculative strategy is positive (negative). Note that the amount of speculative trading the investor executes is capped by the costs of walking the equity book and the terminal liquidation penalty in equity – see the denominator of the speed in (6.16).

Now, we discuss how the volatility of the price of equity affects the speed of trading in equity.

Case 4: High volatility in equity market. When the volatility of P^2 is arbitrarily high,[7] the speed in (6.16) becomes TWAP. To obtain this result, use $\mathbb{E}\left[(P_T^1)^2\right] = V\left[P_T^1\right] + \mathbb{E}\left[P_T^1\right]^2$, where $V[\,\cdot\,]$ denotes the variance operator. Then, as $\mathrm{V}\left[P_T^1\right] \to \infty$, the speed in (6.16) is given by

$$\theta_t^{1,*} = \frac{2\,\alpha_1\, q_0 \left(\mathrm{V}\left[P_T^1\right] + \mathbb{E}\left[P_T^1\right]^2 \right) - \mathbb{E}\left[P_T^2 \left(P_T^1 - \frac{1}{T} \int_0^T P_s^1\, \mathrm{d}s \right) \right]}{2\,\alpha_1\, T \left(\mathrm{V}\left[P_T^1\right] + \mathbb{E}\left[(P_T^1)\right]^2 \right) + 2\, k_1\, \mathbb{E}\left[P_T^2\right]} \to \frac{q_0}{T} , \tag{6.22}$$

which is TWAP.

The intuition behind this result is that as the volatility of the price of equity increases, the chances that the strategy will arrive at time T with a non-negative

[7] We assume that the covariance between P^1 and P^2 remains bounded.

inventory in the stock increase. Thus, it is optimal to employ a strategy that guarantees full liquidation (while minimising the costs of walking the equity LOB) and this is TWAP as discussed in Case 2 above.

Finally, the intuition behind the speed of trading (6.17) in the FX market is the same as that for Case 1 discussed above. When the trend in the currency pair is positive (negative), the investor will buy (sell) the currency pair and close the accumulated position at time T at no cost because $\alpha_2 = 0$.

6.3.4 Example: geometric Brownian motion

Before analysing the performance of the double-execution strategy with financial data from stocks traded in the US and data from the FX market, we discuss some of the features of the strategy with simulated data. To this end, we assume that the dynamics of the midprice of the stock and the mid-exchange rate of the currency pair (that converts the foreign currency into the domestic currency) follow the geometric Brownian motions:

$$\mathrm{d}S_t^1 = \mu_1\, S_t^1\, \mathrm{d}t + \sigma_1\, S_t^1\, \mathrm{d}W_t^1 \qquad \text{and} \qquad \mathrm{d}S_t^2 = \sigma_2\, S_t^2\, \mathrm{d}W_t^2\,, \tag{6.23}$$

respectively, with $S_0^1 = 100$ and $S_0^2 = 1$, and we assume that the spread in both markets is zero.[8] Here, W_t^1, W_t^2 are independent standard Brownian motions, the volatility parameters are $\sigma_1 = 0.007$, $\sigma_2 = 0.002$ over $T = 1$ hour, and the drift parameter of the price of the stock is $\mu_1 \in \{-0.1, -0.05, 0, 0.05, 0.1\}$. Other model parameters are $M = N = 2$, $k_1 = 10^{-3}$, $k_2 = 10^{-4}$, $b_1 = b_2 = 0$, $\alpha_1 = \alpha_2 = 10$, $\phi_1 = \phi_2 = 0$.

The investor must liquidate $q_0 = 1$ shares over a trading window $[0, T]$ with $T = 1$ hour. We train the model with 150 simulated paths of $\boldsymbol{S} = (S_t^1,\, S_t^2)$.[9] The horizon of

[8] The difference between the best bid price and best ask price in the equity LOB is zero; similarly for the FX LOB.

[9] We choose 150 simulations because in our analysis of the performance of the double-execution strategy we employ 147 trading days of Nasdaq and LMAX data to train the model.

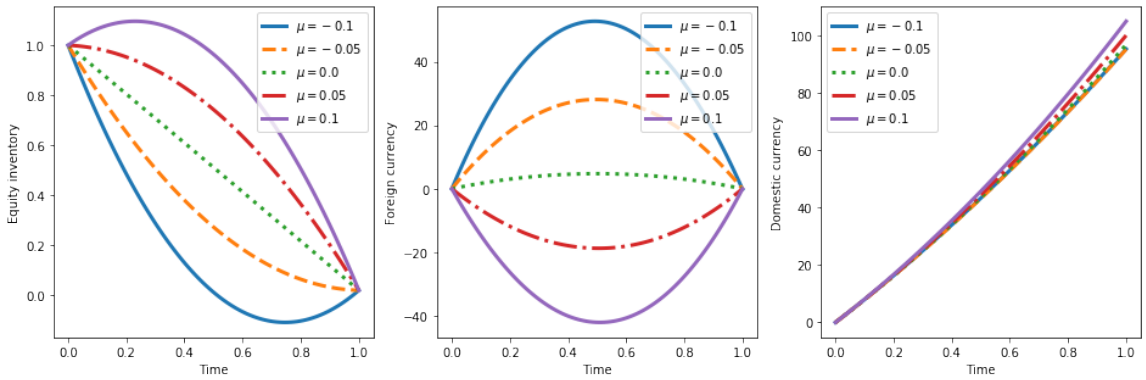

Figure 6.1: Double-execution strategy, averages of 100 simulations: inventory in equity (left panel), inventory in foreign currency (middle panel), inventory in domestic currency (right panel).

each simulation is one hour and the expected signature of the market $\boldsymbol{S}$ is the average of the realised signatures of the simulated paths.

After training the model, we run 100 simulations of $\boldsymbol{S}$ and employ the double-execution strategy developed above, where the pace of trading is 100 market orders at equally spaced time intervals. That is, the investor sends a market order every 0.6 minutes. Figure 6.1 shows the average of the inventory in: equity (left panel), foreign currency (middle panel), domestic currency (right panel).

When $\mu_1 = 0.1$, the left panel of the figure shows that the investor speculates in the equity market to take advantage of an expected increase in the value of the stock. Thus, despite the target is to sell $q_0 = 1$ shares by the terminal time $T = 1$ hour, the investor purchases shares during approximately the first 20 minutes of the execution window, after which the investor starts to sell shares. There are no incentives to speculate in the currency pair because the drift of the mid-exchange rate in the FX market is zero.

The middle panel of Figure 6.1 shows the inventory in the foreign currency. When $\mu_1 = 0.1$ and $t > 0$, the inventory in the foreign currency is short for two reasons. One, the strategy buys shares at the beginning of the trading horizon, so it accumulates a

short position in the foreign currency to finance the speculative purchase of the shares. The pressure to increase the short position in the foreign currency to speculate in the stock market ceases after approximately one third of the trading horizon.

Two, the right panel shows the investor's wealth in the domestic currency, which is increasing for all $t > 0$ and for all values of the drift parameter μ_1. This is optimal because market orders in the FX book incur temporary price impact costs, i.e., walk the LOB, so market orders are spread over the entire trading horizon. In the particular case $\mu_1 = 0.1$, the increase in the inventory of the domestic currency at the beginning of the trading horizon is financed by shorting the foreign currency because during that period the strategy is not selling the stock. Note also that when $\mu_1 = 0.1$, the strategy begins to sell the stock from approximately $t = 20$ minutes onwards. This selling activity increases the position in the foreign currency, but for approximately $t \in [20, 30]$, the short position in the inventory of foreign currency continues to increase (i.e., inventory becomes smaller) because the optimal quantity of domestic currency the investor must purchase cannot be solely financed by the proceeds from selling the stock.

For $\mu_1 = -0.1$, the intuition of the strategy is similar to that discussed above. The investor speculates in the equity market by seeking a short position in the stock to take advantage of the expected fall in its price. That is, the investor sells more shares than the target and closes the short position by the terminal time T.

When the drift of the price of the stock is $\mu_1 = 0$, there are no incentives to speculate in the equity market and Figure 6.1 shows that the execution strategy in the equity market is approximately TWAP. Observe that the inventory path in the foreign currency is concave. At the beginning of the execution window, the speed at which the strategy receives foreign currency is higher than the speed at which the strategy acquires domestic currency.

Finally, observe that the terminal wealth of the strategy is highest (lowest) when the value of the drift parameter μ_1 is highest (lowest). This results from initially holding a long position in the stock that is expected to appreciate (depreciate).

6.4 Data and price impact

We employ LOB data from two sources during the 271 trading days in the period 01/06/2018 to 30/06/2019. The equity data are from Nasdaq and the FX data are from LMAX Exchange. The Nasdaq data consists in the LOB activity – we employ up to the first ten levels of the LOB, which are available at a nanosecond frequency. The FX data consist of the following: volumes posted at the best ask and best bid exchange rates, and all the liquidity taking trades that walked the FX book.

We split the data into a training set and a performance set. The training set runs from 01/06/2018 to 31/12/2018 and we employ it to estimate: the temporary price impact parameters k_1 and k_2, the expected signature of the equity and FX markets, and to derive the coefficients β and γ of the strategy $\boldsymbol{\theta}^*$ that maximises the performance criterion of the investor.

The performance data set runs from 01/01/2019 to 30/06/2019. Below, in Section 6.5, we run the optimal strategy over this period and compare the performance of the double-execution strategy developed above against two benchmarks.

6.4.1 Equity: temporary price impact

In our study, we use data for the companies: Apple, Amazon, FARO Technologies, Tesla, Netflix, International Business Machines Corporation, Boeing Company, American Financial Group, Google, and Microsoft with ticker symbols AAPL, AMZN, FARO, TSLA, NFLX, IBM, BA, AFG, GOOGL, MSFT, respectively. Table 6.1 reports the mean and standard deviation of the daily: (i) number of updates in the

Stock	TOB updates (daily)	Market orders (daily)	*ADV* (in USD)
AAPL	457,503 (122,826)	44,921 (19,748)	786,572,348 (312,751,884)
AMZN	111,491 (42,651)	28,992 (12,932)	1,488,904,272 (618,289,356)
FARO	4,382 (2,281)	270 (122)	851,406 (460,509)
TSLA	96,941 (30,411)	24,443 (11,855)	494,970,622 (257,754,225)
NFLX	119,061 (37,707)	28,888 (13,346)	582,341,525 (344,045,532)
IBM	76,147 (36,783)	6,586 (3,554)	66,375,118 (42,959,626)
BA	57,788 (28,840)	7,507 (3,010)	135,861,318 (54,442,557)
AFG	8,554 (3,492)	747 (314)	4,256,230 (1,804,479)
GOOGL	112,854 (56,140)	14,139 (4,987)	427,571,476 (172,153,972)
MSFT	932,562 (275,408)	34,800 (21,250)	387,008,819 (215,615,544)

Table 6.1: Mean values (standard deviation in parentheses). Training period: 01/06/2018 to 31/12/2018.

best prices (bid and ask) and volumes, which we refer to as top-of-book (TOB), (ii) number of market orders, and (iii) the ADV which is the daily average traded volume in USD.

Recall that we assumed that the temporary price impact of market orders in the equity market is linear in the speed of trading with impact parameter k_1, see (6.2). To estimate the equity parameter k_1, we first compute the cost of walking the LOB; that is, compute the difference between the best bid price available in the LOB when a market sell order is sent and the volume-weighted average price the sell market order receives. We do this for a range of volumes of the sell market order and employ the results to estimate the parameter k_1 as follows.

Take a snapshot of the buy side of the equity LOB every 30 seconds between 11am and 1pm of each day of the training period. For each snapshot of the book, compute the costs of walking the book for market sell orders of size

Stock	$\widehat{k}_1$	Stock	$\widehat{k}_1$
AAPL	1.7326×10^{-5} (1.0652×10^{-5})	IBM	3.2389×10^{-5} (1.8861×10^{-5})
AMZN	4.7309×10^{-4} (1.5782×10^{-4})	BA	1.4361×10^{-4} (7.8140×10^{-5})
FARO	4.6018×10^{-4} (2.0097×10^{-4})	AFG	4.4217×10^{-4} (1.7752×10^{-4})
TSLA	1.4168×10^{-4} (3.7268×10^{-5})	GOOGL	3.7318×10^{-4} (9.4596×10^{-5})
NFLX	1.1631×10^{-4} (3.6605×10^{-5})	MSFT	5.8088×10^{-6} (2.9232×10^{-6})

Table 6.2: Mean of daily estimates (standard deviation of mean in parentheses). Training period: 01/06/2018 to 31/12/2018.

$\{10, 110, 210, 310, 410, \ldots, 1910\}$ shares.

For each day, obtain the parameter estimate $\widehat{k}_1$ from an ordinary least squares regression of the relationship between the costs from walking the book and the volume of the sell market orders. Table 6.2 reports the mean and the standard deviation of the daily estimates of the parameter k_1 for the stocks we study.

6.4.2 FX: temporary price impact

Table 6.3 reports summary statistics for the currency pairs we employ in our study.

We regress the costs of walking the FX LOB on the volume of the buy and sell liquidity taking orders that walked the book during the training period to estimate the impact parameter k_2. The regression is performed twice for each pair of currencies to obtain the price impact parameter in each direction of a trade. Recall that in the model above, the execution price received by the investor who trades in a given direction, say EUR/USD, is $\widetilde{P}_t^{\theta^2} = P_t^{\theta^2} - k_2\,\theta_t^2$, where $P_t^{\theta^2}$ is the bid exchange rate of the currency pair EUR/USD, i.e., the number of units of USD she receives in exchange for EUR. Now, assume that the investor trades in the opposite direction, so she is interested in the bid exchange rate of the 'pair' USD/EUR, i.e. the number of units of EUR she receives in exchange for USD. Then, by no-arbitrage, the bid

Pair	Traded volume (daily, in USD)	Trades (daily)	TOB updates (daily)	Updates TOB (4pm-6pm)
EUR/USD	1,938,083,141 (219,038,378)	32,735 (5,043)	1,147,076 (271,426)	144,075 (34,581)
GBP/USD	883,453,506 (98,683,388)	12,879 (1,677)	1,029,569 (235,099)	129,130 (29,725)
EUR/GBP	126,391,180 (13,748,309)	2,813 (328)	597,482 (135,883)	75,012 (17,040)
USD/JPY	1,430,910,725 (187,102,136)	14,913 (1,910)	867,645 (266,579)	108,828 (33,346)
EUR/CHF	52,941,590 (2,842,850)	839 (344)	462,470 (98,914)	57,938 (12,234)
EUR/MXN	1,142,512 (618,702)	28 (20)	211,667 (41,772)	26,507 (5,207)
GBP/MXN	149,961 (57,488)	23 (18)	329,100 (72,781)	41,229 (9,093)
USD/MXN	72,251,988 (18,646,050)	627 (363)	438,126 (131,341)	54,879 (16,294)
EUR/NZD	16,054,387 (3,434,301)	633 (151)	634,094 (203,569)	79,427 (25,264)
NZD/USD	156,106,521 (20,832,695)	2,057 (506)	419,865 (116,437)	52,560 (14,420)
NZD/CHF	2,705,213 (1,014,065)	316 (87)	125,602 (24,134)	15,793 (3,099)
BTC/USD	731,441 (2,599,774)	45 (105)	544,159 (726,353)	68,450 (90,976)

Table 6.3: Daily mean and standard deviation of volume and number of trades. The time 11:00am to 1:00pm in New York corresponds to 4:00pm to 6:00pm in London, except for the period 28/10/2018 to 04/11/2018, where due to the daylight savings, 11:00am to 1:00pm in New York corresponds to 3:00pm to 5:00pm in London. Training period: 01/06/2018 to 31/12/2018.

exchange rate of the 'pair' USD/EUR corresponds to the inverse of the ask exchange rate of the EUR/USD currency pair, i.e., $1/P_t^{\theta^2}$. In this case, we cannot model the temporary price impact as

$$\frac{1}{P_t^{\theta^2}} - \frac{1}{P_t^{\theta^2} + k_2\,\theta_t^2}\,. \tag{6.24}$$

Thus, we create a synthetic book to trade in the opposite direction and to estimate a temporary price impact that is consistent with our modelling assumptions above.

For example, for the currencies EUR and USD, the convention is to quote the currency pair EUR/USD. Thus, the LOB for EUR/USD quotes the amount of USD that is exchanged for one Euro, so we use the temporary price impact parameter of EUR/USD to indicate the cost of walking the LOB when a market participant aggressively exchanges Euros into USD.

On the other hand, to compute the temporary price impact parameter for an agent who aggressively exchanges USD into Euros we proceed as follows. We use the EUR/USD LOB data to construct a synthetic USD/EUR LOB. We also employ the liquidity taking orders that were sent to the EUR/USD exchange to create synthetic market orders. Specifically, a buy market order in the EUR/USD LOB becomes a sell market order in the synthetic USD/EUR LOB, where the volume of the synthetic market order is also adjusted.[10]

The results are summarised in Table 6.4. On the left-hand side of the table we report the temporary price impact parameter for the LOB of 12 currency pairs, including BTC/USD; recall that BTC denotes the cryptocurrency Bitcoin. The right-

[10]Consider the following example: A trader sends a buy liquidity taking order to the EUR/USD LOB for 100 contracts. Each contract is for 10,000 EUR, thus the trader will purchase 1,000,000 EUR (base currency) and will pay with USD (quote currency). Suppose that the best available price on the ask side of the LOB is 1.0905 USD per unit of EUR, and the average price paid for the transaction after walking the book is 1.1002 USD per unit of EUR. In the synthetic USD/EUR LOB the transaction is as follows: The trader sends a sell liquidity taking order to the synthetic USD/EUR LOB for $100 \times 1.1002 = 110.02$ contracts, where each contract is for 10,000 USD. The best available price in the bid side of the synthetic book is 1/1.0905 EUR per unit of USD, and the average price received for the transaction is 1/1.1002 EUR per unit of USD.

Pair A/B	Temporary price impact A/B	Temporary price impact B/A
EUR/USD	1.5015×10^{-7}	9.7812×10^{-8}
GBP/USD	1.4529×10^{-7}	6.6960×10^{-8}
EUR/GBP	1.8550×10^{-7}	2.6507×10^{-7}
USD/JPY	2.6477×10^{-5}	1.8700×10^{-11}
EUR/CHF	4.7206×10^{-7}	3.0972×10^{-7}
EUR/MXN	8.0689×10^{-6}	6.8794×10^{-10}
GBP/MXN	1.9850×10^{-3}	1.1347×10^{-7}
USD/MXN	5.3701×10^{-6}	7.2553×10^{-10}
EUR/NZD	1.4221×10^{-6}	3.1816×10^{-7}
NZD/USD	1.0640×10^{-7}	3.4428×10^{-7}
NZD/CHF	7.2985×10^{-7}	2.3869×10^{-6}
BTC/USD	3.0182×10^{-1}	1.3683×10^{-5}

Table 6.4: Estimate of price impact parameter k_2. Training period: 01/06/2018 to 31/12/2018. All estimates are statistically significant (p value $< .01$) except for the pairs EUR/MXN, GBP/MXN, MXN/EUR, MXN/GBP.

hand side of the table reports the results for the synthetic LOB. There is a relationship between the order of magnitude of the temporary price impact parameter of an FX LOB and its synthetic LOB. For example, consider the USD/JPY LOB. During the training period, the mean of the mid-exchange rate is 111.74 JPY per unit of USD. If we use the temporary price impact parameter for the USD/JPY LOB reported in Table 6.4, $\widehat{k}_2^{\text{USD/JPY}} = 2.6477 \times 10^{-5}$, then we expect the temporary price impact parameter $\widehat{k}_2^{\text{USD/JPY}}$ of the synthetic JPY/USD LOB to be approximately

$$\left(\frac{1}{111.74} - \frac{1}{111.74 + \widehat{k}^{\text{USD/JPY}}}\right) \times (111.74)^{-1} = 1.8978 \times 10^{-11}\,, \tag{6.25}$$

which is close to the estimate $\widehat{k}_2^{\text{USD/JPY}} = 1.8700 \times 10^{-11}$ we obtained and reported in Table 6.4.

6.5 Performance of double-execution strategy

Every day, the investor liquidates a position in a foreign asset and exchanges the foreign currency into the investor's domestic currency. The trading horizon of the

daily execution programme is between 11:00am and 1:00pm (Eastern time) during the performance period, which is between 01/01/2019 and 30/06/2019. We compare the performance of the signature method developed above with two benchmarks that we describe below in Section 6.5.1.

We assume that the permanent price impact parameters b_i are zero because we employ transaction data from the equity, FX, and cryptocurrency markets and we cannot assume that the trades of the investor affect the data we employ. Also, for simplicity, we assume the investor does not penalise running inventory, i.e., the penalty parameters ϕ_1, ϕ_2 are zero.

We employ the data set of the training period to find the parameters β and γ that maximise the double-execution problem in the form of the polynomial given in (6.12). Note that the coefficients of the polynomial depend on: terms of the expected signature of the market $\mathbf{E}$, time horizon of the liquidation programme T, penalty parameters α_1, α_2, and price impact parameters k_1, k_2.

We employ the data of the training period to estimate the expected signature of the price process $\boldsymbol{P}$ as follows: Compute the realised signature $\widehat{\mathbb{X}}_{0,T}^{<\infty}$ for every day of the training period over the window 11:00am to 1:00pm.[11] Take the average of the realised signatures, which we use as the estimate of the expected signature of the market $\mathbf{E}$ during the performance period.

The other parameters and quantities of the model are as follows. For each stock, the investor liquidates

$$q_0 = HAT \times 2\% \tag{6.26}$$

shares, where HAT is the average hourly traded volume between the hours 11.00am and 1:00pm (Eastern time) and the corresponding GMT hours for the FX market,

[11] We employed the publicly available software `iisignature`, `https://github.com/bottler/iisignature`. Other packages are `esig` and `signatory`, see `https://pypi.org/project/esig/`, and `https://github.com/patrick-kidger/signatory`.

Stock	**Sell MOs**	**Buy MOs**	Sell ***AHV***	Buy ***AHV***	Sell ***HAT***	Buy ***HAT***
AAPL	6,446 (3,221)	4,704 (2,087)	105,610,475 (50,160,550)	81,971,518 (36,393,728)	535,588 (271,552)	415,011 (193,943)
AMZN	5,443 (2,895)	1,853 (925)	236,786,112 (121,550,656)	117,409,895 (58,886,123)	136,567 (75,440)	67,906 (37,479)
FARO	38 (24)	23 (15)	113,744 (83,449)	84,942 (82,661)	2,066 (1,486)	1,497 (1,345)
TSLA	4,069 (2,340)	2,055 (1,096)	77,901,768 (46,853,114)	45,905,201 (25,777,826)	244,002 (145,126)	144,129 (80,256)
NFLX	5,151 (3,005)	2,121 (1,222)	92,949,030 (71,408,427)	51,034,595 (39,243,363)	276,741 (197,859)	151,878 (107,575)
IBM	993 (540)	711 (358)	9,865,777 (6,491,904)	6,852,669 (4,263,315)	73,477 (52,652)	51,082 (34,876)
BA	1,221 (617)	703 (298)	21,734,486 (11,483,325)	13,358,209 (5,531,983)	62,435 (32,840)	38,431 (15,972)
AFG	109 (56)	77 (41)	595,083 (323,190)	400,153 (232,665)	5,704 (3,328)	3,820 (2,291)
GOOGL	2,312 (960)	1,021 (469)	62,871,701 (29,440,629)	32,663,815 (17,199,660)	55,598 (28,189)	28,910 (16,783)
MSFT	4,903 (3,158)	3,835 (2,669)	52,726,082 (32,850,013)	39,150,944 (25,405,701)	497,026 (317,840)	368,935 (245,088)

Table 6.5: Mean and standard deviation in parentheses. Data from 11:00am to 1:00pm Eastern time in the training period from 01/06/2018 to 31/12/2018. HAT denotes the hourly average number of shares traded between 11:00am and 1:00pm.

see Table 6.1.

The terminal liquidation penalty parameters are $\alpha_1 = \alpha_2 = 100$, and the trading horizon is 2 hours (i.e., from 11:00am to 1:00pm Eastern time).

6.5.1 Benchmarks

We compare the terminal wealth that the double-execution strategy obtains each day of the performance period against two benchmarks. The benchmarks are based on the TWAP strategy, which is a constant execution strategy. TWAP is widely used in execution algorithms and as a benchmark to compare the performance of trading strategies, see [26] and [70]. For example, assume a trader wishes to employ n market orders to sell q_0 shares over the interval $[0, T]$. In discrete-time, the TWAP strategy sends a market sell order every T/n units of time for the quantity q_0/n.[12] In this

[12]In continuous-time, the TWAP strategy executes at the speed q_0/T.

chapter we employ two benchmarks: the Double-TWAP and the TWAP-Aggressive.

Double-TWAP. The investor executes all the inventory in the foreign stock with a TWAP strategy over the period $[0, T/2]$, after which the investor exchanges the foreign currency (i.e., USD) into the domestic currency with a TWAP strategy over the period $(T/2, T]$.

TWAP-Aggressive. The investor employs TWAP over the period $[0, T]$ to liquidate the inventory in the foreign stock. Every time the investor collects USD from selling a batch of the equity inventory, she executes a market order in the FX market to exchange the entire position in USD into the domestic currency.

6.5.2 Frequency of trading of double-execution strategy

As discussed above, we cannot incorporate the effect that the investor's trading activity would have on the transaction data in the equity and FX markets. Nevertheless, we choose the investor's trading pace similar to that usually seen in the markets so that the investor's activity 'camouflages' within the usual order flow in the equity and FX markets.

In the equity market, the investor sends liquidity taking orders at a rate similar to what the market usually bears during the investor's execution window. Specifically, in the equity market, we assume that the number of market orders the investor sends is equal to the average number of sell market orders for the corresponding stock over the same trading window in Nasdaq during the training period.

Observe that the mean trading activity for the stocks in Nasdaq varies considerably among the stocks we study. For example, the mean number of liquidity taking sell orders for AAPL is nearly six times that for BA, so the investor will send more orders to the AAPL book than to the BA book over the trading window – recall that the liquidation target is a percentage of the average volume traded over those two hours,

see (6.26) above. Also, observe that the standard deviation of the trading activity for all stocks is also very high, which makes it easier to camouflage the additional market sell orders that the investor sends to the equity market. Finally, we note that Table 6.5 only includes data for Nasdaq, which represents a share of the equity market between 15% and 20%.[13]

We assume that in the FX market the investor sends the same number of liquidity taking orders as that in the equity market. For example, if we assume that the position to liquidate in AAPL is 2% of the average number of shares sold between 11:00am to 1:00pm, then the investor liquidates 535,588 shares with 6,446 sell market orders over the two-hour trading window. Also, the investor sends 6,446 liquidity taking orders in the FX exchange. Finally, we assume that the orders in both markets are sent at equally spaced points in time, so for AAPL the investor sends market orders in the equity and FX markets every 1.11 seconds.

6.5.3 Performance

For each day, we compare the cash obtained in the domestic currency with the double-execution strategy to that obtained with the Double-TWAP and with the TWAP-Aggressive. Specifically, we compute the following performance measure in basis points (bps):

$$\frac{\mathcal{C}^{\boldsymbol{\theta}^*} - \mathcal{C}^{\boldsymbol{\theta}^B}}{\mathcal{C}^{\boldsymbol{\theta}^B}} \times 10^4 \,. \tag{6.27}$$

Here, $\mathcal{C}^{\boldsymbol{\theta}^*}$ is the cash in domestic currency obtained by the double-execution strategy and $\mathcal{C}^{\boldsymbol{\theta}^B}$ is the cash in the domestic curency obtained by the benchmark B, where B denotes the Double-TWAP or the TWAP-Aggressive.

[13] See `www.nasdaqtrader.com`.

	Benchmark: Double-TWAP				
	Domestic Currency				
	GBP	**EUR**	**JPY**	**MXN**	**NZD**
AAPL	19.90 (7.96)	22.68 (12.67)	26.91 (17.51)	32.22 (18.09)	38.70 (17.85)
AMZN	21.38 (12.40)	25.21 (10.82)	21.79 (15.96)	26.42 (16.34)	22.58 (18.88)
FARO	299.36 (43.18)	299.84 (42.82)	399.79 (32.21)	402.34 (31.71)	402.15 (32.71)
TSLA	27.22 (9.09)	26.51 (29.49)	30.78 (19.62)	34.82 (20.98)	31.99 (18.75)
NFLX	25.81 (13.64)	26.21 (13.11)	30.31 (11.65)	27.86 (13.83)	31.70 (23.13)
IBM	26.69 (14.10)	23.60 (11.63)	28.99 (12.86)	27.81 (13.65)	29.98 (13.40)
BA	21.44 (18.91)	25.37 (10.41)	19.50 (10.20)	26.22 (20.48)	26.20 (20.22)
AFG	178.54 (44.18)	174.95 (9.24)	177.48 (9.21)	177.78 (9.89)	175.57 (9.20)
GOOGL	19.24 (12.20)	22.44 (15.40)	23.33 (15.71)	24.35 (16.22)	23.20 (15.72)
MSFT	16.01 (13.11)	17.92 (13.71)	34.40 (13.33)	16.39 (14.23)	41.94 (13.57)

Table 6.6: Mean and standard deviation of performance measure (6.27). Performance period: 01/01/2019 to 30/06/2019 and trading horizon 11:00am to 1:00pm (Eastern time).

6.5.3.1 Equity and Foreign Exchange Trading

Table 6.6 shows the mean and standard deviation of the performance measure (6.27) for various stocks and for various domestic currencies. All stocks are traded in USD in Nasdaq and we assume that the domestic currency of the investor is one of the following: GBP, EUR, JPY, MXN, NZD.

In all stocks and domestic currencies, the optimal signature strategy outperforms the Double-TWAP benchmark, see Table 6.6. The outperformance is highest for FARO. Also, the signature trading strategy outperformed the Double-TWAP between 95% and 99% of the days when executing IBM, while it outperformed the benchmark on 100% of the days when executing FARO, see Table 6.7. Furthermore, in all the cases we report here, the trading speeds always remained positive, i.e., the investor

	Benchmark: Double-TWAP				
	Domestic Currency				
	GBP	**EUR**	**JPY**	**MXN**	**NZD**
AAPL	86.29%	92.74%	95.16%	91.12%	95.96%
AMZN	93.54%	91.93%	93.52%	95.96%	89.79%
FARO	100.00%	100%	100.00%	100.00%	100.00%
TSLA	86.80%	85.48%	86.29%	85.48%	87.90%
NFLX	88.70%	88.71%	90.32%	88.71%	89.51%
IBM	95.16%	96.77%	99.19%	97.58%	99.20%
BA	88.70%	91.12%	82.25%	91.93%	93.54%
AFG	100%	100%	100%	100%	100%
GOOGL	88.93%	92.74%	91.93%	91.11%	91.93%
MSFT	90.38%	91.34%	99.03%	88.46%	100%

Table 6.7: Percentage of days when the signature trading strategy outperforms the Double-TWAP benchmark. Performance period: 01/01/2019 to 30/06/2019 and trading horizon 11:00am to 1:00pm (Eastern time).

never executes speculative trades in the equity or FX market.

Figure 6.2 shows the average (during the performance period) of: inventories of equity, foreign currency (USD), and domestic currency (MXN), for the double-execution strategy and the two benchmarks. Observe that at the beginning (end) of the trading window the double-execution strategy liquidates the inventory in equity faster (slower) than the TWAP-Aggressive strategy. Also, the right panel shows that the wealth in the domestic currency of the double-execution strategy is most of the time below that of the TWAP-Aggressive. However, by the terminal date we know that the double-execution strategy is the best as shown in Tables 6.6 and 6.8. Note also that the double execution-strategy outperforms the TWAP-Aggressive in every day of the performance period for FARO and AFG, see Table 6.7.

6.5.4 Augmenting information of signatures

We augment the market by including a third asset, so $\boldsymbol{P}^{\boldsymbol{a}} = (P_t^1, P_t^2, P_t^3)$, where P_t^3 denotes the prices of the SPDR S&P 500 with ticker SPY.[14] In our setup, the investor

[14]See `https://www.nasdaq.com/glossary/s/spdr`.

	Benchmark: TWAP-Aggressive				
	Domestic Currency				
	GBP	**EUR**	**JPY**	**MXN**	**NZD**
AAPL	86.49 (12.70)	83.77 (1.84)	71.26 (1.52)	108.76 (1.92)	92.04 (1.93)
AMZN	9.18 (0.56)	13.00 (5.38)	8.74 (0.56)	27.02 (5.69)	26.76 (3.01)
FARO	217.94 (1.97)	218.53 (1.97)	317.19 (10.74)	319.25 (11.49)	319.25 (10.44)
TSLA	34.74 (3.25)	34.03 (2.68)	30.56 (2.55)	33.80 (2.61)	29.19 (2.72)
NFLX	30.66 (0.71)	30.57 (0.49)	25.99 (0.32)	29.61 (0.73)	24.54 (0.52)
IBM	43.16 (18.44)	109.33 (1.24)	109.21 (1.34)	111.83 (1.30)	108.22 (1.57)
BA	23.30 (20.84)	27.19 (0.88)	32.22 (0.63)	26.77 (1.07)	25.35 (0.77)
AFG	72.94 (37.11)	69.43 (6.01)	71.33 (7.39)	71.26 (7.23)	69.06 (6.82)
GOOGL	46.16 (28.82)	36.57 (4.24)	65.38 (1.74)	11.20 (4.15)	10.10 (3.97)
MSFT	11.30 (5.99)	15.42 (4.68)	10.61 (3.78)	18.50 (1.65)	157.83 (1.76)

Table 6.8: Mean and standard deviation of performance measure (6.27). Performance period: 01/01/2019 to 30/06/2019 and trading horizon 11:00am to 1:00pm (Eastern time).

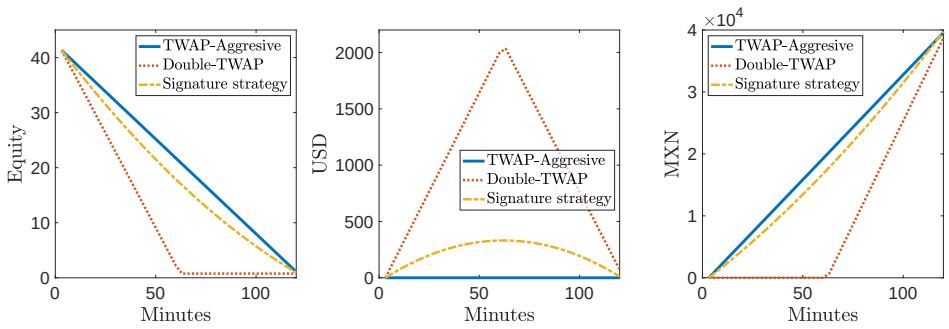

Figure 6.2: Equity, foreign currency and domestic currency inventory for the signature method, Double-TWAP and TWAP-Aggressive strategies for the execution of FARO into MXN.

	Benchmark: Double-TWAP	
	Domestic Currency	
	GBP	GBP
	Market: $\boldsymbol{P} = (P_t^1, P_t^2)$	Market: $\boldsymbol{P^a} = (P_t^1, P_t^2, P_t^3)$
AAPL	19.90 (7.96)	22.93 (18.20)
AMZN	21.38 (12.40)	25.31 (18.49)
FARO	299.36 (43.18)	396.06 (44.52)
TSLA	27.22 (9.09)	30.05 (10.53)
NFLX	25.81 (13.64)	31.25 (22.14)
IBM	26.69 (14.10)	37.15 (14.64)
BA	21.44 (18.91)	24.92 (21.44)
AFG	178.50 (44.18)	201.74 (12.86)
GOOGL	19.24 (12.20)	28.26 (43.30)
MSFT	16.01 (13.11)	17.77 (14.45)

Table 6.9: Mean and standard deviation of performance measure (6.27) with additional information provided by SPY (third column). Performance period: 01/01/2019 to 30/06/2019 and trading horizon 11:00am to 1:00pm (Eastern time). The second column are the results of the double-execution strategy without SPY.

follows the double-execution strategy with the additional information provided by SPY. As above, we employ the training period to compute the expected signature of the market $\boldsymbol{P^a}$ and to obtain the optimal signature strategy parameters β, γ.

Table 6.9 shows the mean and standard deviation of the performance measure (6.27) for various stocks and for the domestic currency GBP. For comparison purposes, the second column reports the same results as those in Table 6.6 where the market consisted of equity in Nasdaq and the GBP/USD currency pair. The last column of Table 6.9 shows the results when we include the SPY. It is clear that the additional information provided by the dynamics of SPY improves the performance of the double-execution strategy.

	Benchmark: Double-TWAP				
	Domestic Currency				
	GBP	**EUR**	**JPY**	**MXN**	**NZD**
BTC/USD	261.26 (0.98)	194.24 (0.68)	290.97 (0.48)	194.39 (0.38)	286.43 (0.56)

Table 6.10: Mean and standard deviation of performance measure (6.27). Performance period: 01/01/2019 to 30/06/2019 and trading horizon 11:00am to 1:00pm (Eastern time).

6.5.5 Illiquid trading

6.5.5.1 Cryptocurrency Trading

As of March 2020, market makers in LMAX Digital[15] offer liquidity in currency pairs to exchange main cryptocurrencies into the currencies EUR, JPY, USD. For example, when an agent wishes to exchange a large position in BTC into one of the following: GBP, MXN, NZD, there might be no currency pair to exchange BTC directly into the desired currrency. That is, either the market does not trade the pairs BTC/GBP, BTC/MXN, BTC/NZD or there is no liquidity to meet the investor's demand. Therefore, the investor employs a triplet of currency pairs to exchange BTC into the desired domestic currency via an intermediate currency pair. The solution to this optimal trading problem is that given by the double-execution strategy derived above.

We follow the approach in Section 6.5 to illustrate the performance of the double-execution strategy to trade triplets of currency pairs. The investor employs the currency USD as the link between BTC and the domestic currency. The investor holds a position in BTC equivalent to 0.01% of the ADV of the pair BTC/USD. Table 6.11 summarises the results of the performance measure (6.27) with the benchmarks Double-TWAP and TWAP-Aggressive.

[15]See `https://www.lmaxdigital.com/`.

	Benchmark: TWAP-Aggressive				
	Domestic Currency				
	GBP	**EUR**	**JPY**	**MXN**	**NZD**
BTC/USD	555.44 (7.58)	562.21 (8.34)	523.65 (7.55)	562.21 (8.58)	524.67 (6.93)

Table 6.11: Mean and standard deviation of performance measure (6.27). Performance period: 01/01/2019 to 30/06/2019 and trading horizon 11:00am to 1:00pm (Eastern time).

6.5.5.2 Illiquid Currency Pairs

Here, we analyse how the investor circumvents costs that stem from trading in illiquid currency pairs. We build on Subsection 6.5.5.1 and consider the case in which the agent wishes to exchange a large position in NZD into CHF. The currency pair NZD/CHF is illiquid, whereas the currency pairs USD/NZD and USD/CHF are liquid. We compare the following two strategies: (i) the double-execution strategy uses NZD/USD and USD/CHF to exchange a position of NZD into CHF, and (ii) a TWAP strategy in NZD/CHF to transfer a position of NZD into CHF. See [25] for a model where the agent liquidates a large position over an illiquid currency pair by means of a triplet of currencies.

The initial position of the investor in NZD is 0.01% of the ADV. We compare (i) and (ii), and we find that the outperformance of the double-execution strategy in NZD/USD and USD/CHF when compared with the TWAP strategy in NZD/CHF is 19.22 basis points with a standard deviation of 6.12.

6.6 Signatures, neural networks, and optimal control

Recall that an important step in our approach is to approximate the speed of trading in the equity and FX markets with signature trading strategies. This step enables us

to write the investor's double-execution problem as a polynomial that is straightforward both to train and to maximise to derive the optimal signature trading strategy.

As discussed above, an alternative approach is to solve the double-execution problem with NN-based trading strategies, where the input at time $t \in [0, T]$ is the realised path of $\boldsymbol{P}$ up to time t, and the output is the amount of equity and the currency pair the investor needs to trade. This approach is feasible if the investor's performance criterion is differentiable with respect to the weights of the NN. The NN can be trained in various ways, one of which is the gradient descent method; see for example [84], [13, 113, 28, 73].

A key difference between the two approaches, signature and NN, is the number of sample paths that are required to train each model and to compute the optimal trading strategy. Compared with signature-based approaches, NNs need a large training data set. For example, [28] generate 15,000 paths to train a NN to solve an investment problem. The work of [13] employs approximately 10^7 sample paths of daily prices for a few months up to a year to hedge financial derivatives.

If we were to employ NNs to solve the double-execution problem discussed in this chapter, it would require much more data than the 147 equity and FX paths (i.e., the training data set) we employ to train the investor's signature-based model. The signature method requires much fewer price paths to train the model because it uses the training data only to estimate the expected signature of the market. Clearly, if the size of the training data set increases, the estimate of the expected signature of the market improves.

Another difference between both approaches is the method one employs to find the optimal trading strategies. Methods based on gradient descent, which are commonly used to train NNs, are in many cases notoriously slow to obtain a (local) minimum. In contrast, one obtains the optimal signature-based strategies very quickly because

the signature method reduces the investor's problem to finding the extrema of a polynomial such as the one in (6.12) – for example, it takes approximately one minute to solve the optimisation problem for the example in Section 6.3.4.

On the other hand, an important advantage of the NN approach is that it can be employed in a wide range of problems and model setups. In general, if the investor's performance criterion is differentiable with respect to the network's weights, the NN approach is useful provided the training data set is large enough.

The signature method is employed in a range of financial problems (e.g., [86], [101]), and it relies on expressing functions as signature-based polynomials, which is not always feasible. When one cannot express the investor's performance criterion as a polynomial, one alternative is to employ the gradient descent method to obtain optimal signature trading strategies, see [87]. As with NNs, the success of this alternative route relies on training the model on a large data set, which comes at the cost of slow convergence in the training phase and in the derivation of the optimal strategies.

Finally, recall that in Section 6.3.2 we mentioned that the investor can increase the dimension of $\boldsymbol{P}$ with other features (traded volume, prices of other assets, order flow, etc) to augment the market information employed by the investor in the double-execution problem. However, there is a tradeoff between enhancing market information and the computational tractability of the investor's problem because the dimension of the signature trading strategies increases as the dimension of $\boldsymbol{P}$ increases. Specifically, when a number of n new features are added to $\boldsymbol{P}$, the dimension of the strategies grows by an amount in the order of $\mathcal{O}(n^M + n^N)$, where M, N are parameters in Definition 6.3.1. We note that NNs also suffer from a similar explosion in dimensionality when more features are included in the model.

6.7 Conclusions

We solved a double-execution problem where an investor liquidates a large position in a foreign currency and exchanges the proceeds of the sale into the investor's domestic currency. We employed signature-based techniques to reduce the investor's problem to a polynomial that is straightforward both to train with financial data and form which it is easy to obtain the double-execution strategies.

We employed high-frequency transaction data from Nasdaq and from LMAX Exchange. Our results showed that the double-execution strategy considerably outperforms two benchmarks that are based on TWAP execution strategies. We also employed the double-execution strategy to trade cryptocurrencies and illiquid currency pairs.

Directions of future research based on signatures include the following: Our approach can be extended to solve other investment problems. For example, investors who need to rebalance portfolios or those who must employ domestic currency to purchase a large position of a security (or securities) in a foreign financial market.

Part III

Signatures in machine learning

Chapter 7

Deep Signature Transforms

The signature is an infinite graded sequence of statistics known to characterise a stream of data up to a negligible equivalence class. It is a transform which has previously been treated in machine learning as a fixed feature transformation, on top of which a model may be built. We propose a novel approach which combines the advantages of the signature transform with modern deep learning frameworks. By learning an augmentation of the stream prior to the signature transform, the terms of the signature may be selected in a data-dependent way. More generally, we describe how the signature transform may be used as a layer anywhere within a neural network. In this context it may be interpreted as an activation function not operating element-wise. We present the results of empirical experiments to back up the theoretical justification. The code is available at `github.com/patrick-kidger/Deep-Signatures`. This chapter, which is based on the paper [87] published at NeurIPS 2019, is joint work with Patric Bonnier, Patrick Kidger, Cristopher Salvi and Terry Lyons.

7.1 Introduction

When data is ordered sequentially then it comes with a natural path-like structure: the data may be thought of as a discretisation of a path $X\colon [0,1] \to V$, where V is some Banach space. In practice we shall always take $V = \mathbb{R}^d$ for some $d \in \mathbb{N}$. For

example the changing air pressure at a particular location may be thought of as a path in $\mathbb{R}$; the motion of a pen on paper may be thought of as a path in $\mathbb{R}^2$; the changes within financial markets may be thought of as a path in $\mathbb{R}^d$, with d potentially very large.

We refer the reader to [32] for a primer on the use of the signature in the context of machine learning. In short, the signature of a path determines the path essentially uniquely, and does so in an efficient, computable way. Furthermore, the signature is rich enough that every continuous function of the path may be approximated arbitrarily well by a linear function of its signature; it may be thought of as a 'universal nonlinearity'. Taken together these properties make the signature an attractive tool for machine learning. The most simple way to use the signature is as feature transformation, as it may often be simpler to learn a function of the signature than of the original path. Over the last decade signatures have been used in various ways in machine learning [137, 135, 140, 138, 97, 141, 139, 90, 34].

The signature is an infinite sequence, so in practice some finite collection of terms must be selected. Since the magnitude of the terms exhibit factorial decay, it is usual [100] to simply choose the first N terms of this sequence, which will typically be the largest terms. These first N terms are called the *truncated signature*, denoted Sig^N. But if the function to be learned depended nontrivially on the higher degree terms, then crucial information has nonetheless been lost.

This may be remedied. Apply a pointwise augmentation to the original stream of data before taking the signature. Then the first N terms of the signature may better encode the necessary information [90, 34]. Explicitly, let $\Phi\colon \mathbb{R}^d \to \mathbb{R}^e$ be fixed; one could ensure that information is not lost by taking $\Phi(x) = (x, \varphi(x))$ for some φ. Then rather than taking the signature of $\mathbf{x} = (x_1, \ldots, x_n)$, where $x_i \in \mathbb{R}^d$, instead take the signature of $\Phi(\mathbf{x}) = (\Phi(x_1), \ldots, \Phi(x_n))$. In this way one may capture higher

order information from the stream in the lower degree terms of the signature.

7.1.1 Our work

But how should this augmentation Φ be chosen? Previous work has fixed it arbitrarily, or experimented with several options before choosing one [90, 34]. Observe that in each case the map $\mathbf{x} \mapsto \mathrm{Sig}^N(\Phi(\mathbf{x}))$ is still ultimately just a feature transformation on top of which a model is built. Our more general approach is to allow the selection of Φ to be data-dependent, by having it be learned; in particular it may be a neural network. Furthermore there is no reason it should necessarily operate pointwise, nor (since it is now learned) need it be of the form $(x, \varphi(x))$. In this way we may enjoy the benefits of using signatures while avoiding their main limitation.

But this means that the signature transform is essentially operating as a layer within a neural network. It consumes a tensor of shape $(b, d, ?)$ – corresponding to a batch of size b of paths in $\mathbb{R}^d$ that have been sampled at different number of times – and returns a tensor of shape $(b, (d^{N+1} - 1)/(d - 1))$, where N is the number of terms used in the truncated signature.[1] The signature is essentially an activation function that does not operate element-wise.

There is no reason to stop here. If the signature layer works well once then it is natural to seek to use it again. The obvious problem is that the signature transform consumes a stream of data and returns statistics which have no obvious stream-like qualities. The solution is to lift the input stream to a *stream of streams*; for example, the stream of data $\mathbf{x} = (x_1, \ldots, x_n)$ may be lifted to the 'expanding windows' of $(\mathbf{x}_1, \ldots, \mathbf{x}_n)$, where $\mathbf{x}_i = (x_1, \ldots, x_i)$. Now apply the signature to each stream to obtain a stream of signatures $(\mathrm{Sig}^N(\mathbf{x}_1), \ldots, \mathrm{Sig}^N(\mathbf{x}_n))$, which is essentially a stream in Euclidean space. And now this new stream may be augmented via a neural network

[1] As $(d^{N+1} - 1)/(d - 1) = \sum_{k=0}^{N} d^k$ is the number of scalar values in a signature with N terms.

and the process repeated again, as many times as we wish.

In this way the signature transform has been elevated from a one-time feature transformation to a first-class layer within a neural network. Thus we may reap the benefits of both the signature transform, with its strong corpus of mathematical theory, and the benefits of neural networks, with their great empirical success. We refer to this usage as *deep signatures*.

The remainder of the paper is laid out as follows. In Section 7.2 we briefly discuss some related work, in Section 7.3 we detail the specifics of embedding the signature as a layer within a neural network. Sections 7.4 and onwards are dedicated to experiments; we demonstrate positive results for generative, supervised, and reinforcement learning problems. Appendix E specifies implementation details.

7.2 Related Work

Some related work is by necessity already discussed in the introduction. We expand a little on their proposed models here.

Definition 7.2.1. *Given a set V, the space of streams of data in V is defined as*

$$\mathcal{S}(V) = \{\mathbf{x} = (x_1, \ldots, x_n) : x_i \in V, n \in \mathbb{N}\}.$$

Given $\mathbf{x} = (x_1, \ldots, x_n) \in \mathcal{S}(\mathcal{X})$, the integer n is called the length of $\mathbf{x}$.

We now define the *signature map* Sig, which acts on streams of data in $\mathbb{R}^d$.

Definition 7.2.2. *Let $\mathbf{x} = (x_1, \ldots, x_n) \in \mathcal{S}(\mathbb{R}^d)$ be a stream of data. Let $X : [0, 1] \to \mathbb{R}^d$ be a linear interpolation of $\mathbf{x}$. The signature map $\mathrm{Sig} : \mathcal{S}(\mathbb{R}^d) \to T((\mathbb{R}^d))$ is defined by*

$$\mathrm{Sig}(\mathbf{x}) := \mathbb{X}_{0,1}^{<\infty} \in T((\mathbb{R}^d))$$

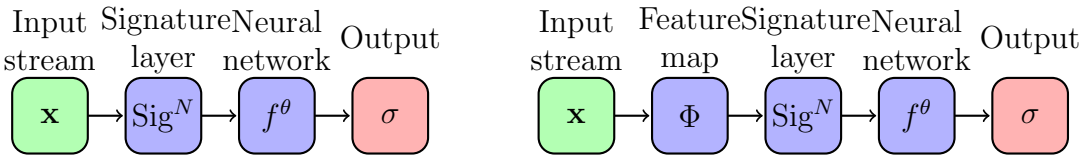

(a) Neural-signature model. Trainable parameters: θ.

(b) Neural-signature-augment model. Trainable parameters: θ.

Figure 7.1: Two simple architectures with a signature layer.

where $\mathbb{X}_{0,1}^{<\infty}$ *is the signature of* X *(Definition 2.1.1). Similarly, we define the truncated signature map* $\mathrm{Sig}^N : \mathcal{S}(\mathbb{R}^d) \to T^{(N)}(\mathbb{R}^d)$ *by*

$$\mathrm{Sig}^N(\mathbf{x}) := \mathbb{X}_{0,1}^{\leq N} \in T^{(N)}(\mathbb{R}^d).$$

Remark 7.2.3. The definition of the signature map may appear to depend on the choice of linear interpolation, as the choice of linear interpolation is not unique. However, by the invariance of signatures under reparametrisation (see Proposition 2.1.3) the signature map as introduced above is well-defined as it does not depend on the choice of linear interpolation.

Remark 7.2.4. By Chen's identity (see Proposition 2.1.4) computing the signature of an incoming stream of data is efficient. Indeed, suppose one has obtained a stream of data $\mathbf{x} \in \mathcal{S}(\mathbb{R}^d)$ and computed its signature. Suppose that after some time more data has arrived, $\mathbf{y} \in \mathcal{S}(\mathbb{R}^d)$. In order to compute the signature of the whole signal one only needs to compute the signature of the new piece of information, and tensor product it with the already-computed signature.

Two simple models utilising the signature layer are shown in Figure 7.1. In principle the universal nonlinearity property of signatures guarantees that the model shown in Figure 7.1a, with the neural network taken to be a single linear layer and the input stream assumed to already be time-augmented, is rich enough to learn any continuous function. In practice, of course, the signature must be truncated. Furthermore, it

is not clear how to appropriately choose the truncation hyperparameter N. Thus a more practical approach is to remove the restriction that the neural network must be linear, and learn a nonlinear function instead. This approach has been applied successfully in various tasks [99, 137, 135, 140, 138, 97, 141, 139].

An alternate model is shown in Figure 7.1b. Following [90, 34], one could also apply a pointwise transformation to the stream before the signature, lifting the d-dimensional stream of data $(x_1, \ldots, x_n) \in \mathcal{S}(\mathbb{R}^d)$ into a higher-dimensional feature space via a feature map $\Phi\colon \mathbb{R}^d \to \mathbb{R}^e$, yielding $(\Phi(x_1), \ldots, \Phi(x_n)) \in \mathcal{S}(\mathbb{R}^e)$. Then the signature of $\Phi(\mathbf{x})$ may potentially capture properties of the stream of data that will yield more effective models.

7.3 The signature as a layer in a neural network

However, there is not always a clear candidate for the feature map Φ and a good choice is likely to be data-dependent. Thus we propose to make Φ learnable by taking $\Phi = \Phi^\theta$ to be a neural network with trainable parameters θ. In this case, we again obtain the neural network shown in Figure 7.1b, except that Φ is now also learnable.

The signature has now become a layer within a neural network. It consumes a tensor of shape $(b, d, ?)$ – corresponding to a batch of size b of paths in $\mathbb{R}^d$ that have been sampled at different number of points – and returns a tensor of shape $(b, (d^{N+1}-1)/(d-1))$, where N is the number of terms used in the truncated signature.

7.3.1 Stream-preserving neural networks

Let $\mathbf{x} = (x_1, \ldots, x_n) \in \mathcal{S}(\mathbb{R}^d)$. Whatever the choice of Φ^θ, it must preserve the stream-like nature of the data if we are to take a signature afterwards. The simplest way of doing this is to have Φ^θ map $\mathbb{R}^d \to \mathbb{R}^e$, so that it operates pointwise. This

defines $\Phi(\mathbf{x})$ by

$$\Phi(\mathbf{x}) = (\Phi^\theta(x_1), \dots, \Phi^\theta(x_n)) \in \mathcal{S}(\mathbb{R}^e). \tag{7.1}$$

Another way to preserve the stream-like nature is to sweep a one dimensional convolution along the stream; more generally one could sweep a whole feedforward network along the stream. For some $m \in \mathbb{N}$ and $\Phi^\theta \colon \mathbb{R}^{d \times m} \to \mathbb{R}^e$ this defines $\Phi(\mathbf{x})$ by

$$\Phi(\mathbf{x}) = (\Phi^\theta(x_1, \dots, x_m), \dots, \Phi^\theta(x_{n-m+1}, \dots, x_n)) \in \mathcal{S}(\mathbb{R}^e). \tag{7.2}$$

More generally still the network could be recurrent, by having memory. Let $\Phi_0 = 0$, fix $m \in \mathbb{N}$, and define $\Phi_k = \Phi^\theta(x_k, \dots, x_{k+m}; \Phi_{k-1})$ for $k = 1, \dots, n - m + 1$. Then define $\Phi(\mathbf{x})$ by

$$\Phi(\mathbf{x}) = (\Phi_1, \dots, \Phi_{n-m+1}) \in \mathcal{S}(\mathbb{R}^e). \tag{7.3}$$

It is worth taking a moment to think what we really mean by 'stream-like nature'. The signature transform is defined on paths; it is applied to a stream of data in $\mathcal{S}(\mathbb{R}^d)$ by first interpolating the data into a path and then taking the signature. The data is treated as a discretisation or set of observations of some underlying path. Note that there is nothing wrong with the path itself having a discrete structure to it; for example a sentence.

In principle one could reshape a tensor of shape (b, nd) with no stream-like nature into one of shape (b, d, n), and then take the signature. However it is not clear what this means mathematically. There is no underlying path. The signature is at this point an essentially arbitrary transformation, without the mathematical guarantees normally associated with it.

7.3.2 Backpropagation

Despite being formed of integrals, the signature is in fact straightforward and efficient to compute exactly. More than that, the computation may in fact be described in

terms of standard tensor operations. As such it may be backpropagated through without difficulty. In particular this means that having a learnable map before the signature does not pose a problem.

However, viewing the signature as a function of the values of the path is, as rough path theory has shown, not a good idea in general. Therefore, an appropriate notion of backpropagation is needed. In other words, how do we perturb the input path when applying backpropagation? In general, we may need to perturb not only the values of the path, but the Lévy area as well, or other higher order terms of the signature. The input path should in general be seen as a rough path, and as such further research is needed to investigate the appropriate notion of backpropagating through signatures.

7.3.3 Multiple signature layers

Applying the signature transform consumes the stream-like nature of the data, which if done naively prevents using the signature layer multiple times. The straightforward solution is to take a stream of signatures in the following way: given a stream $\mathbf{x} = (x_1, \ldots, x_n) \in \mathcal{S}(\mathbb{R}^d)$, let $\mathbf{x}_k = (x_1, \ldots, x_k)$ for $k = 2, \ldots, n$, and apply the signature to each $\mathbf{x}_k$ to obtain the stream

$$(\mathrm{Sig}^N(\mathbf{x}_1), \ldots, \mathrm{Sig}^N(\mathbf{x}_n)) \in \mathcal{S}(\mathbb{R}^{(d^{N+1}-1)/(d-1)}).$$

This notion may be generalised: let

$$\ell = (\ell^1, \ell^2, \ldots, \ell^v) \colon \mathcal{S}(\mathbb{R}^d) \to \mathcal{S}(\mathcal{S}(\mathbb{R}^e)),$$

which we refer to as a lift into the space of streams of streams.[2] Then we apply the signature stream-wise to define $\mathrm{Sig}^N(\ell(\mathbf{x}))$ by

$$\mathrm{Sig}^N(\ell(\mathbf{x})) = \big(\mathrm{Sig}^N(\ell^1(\mathbf{x})), \ldots, \mathrm{Sig}^N(\ell^v(\mathbf{x}))\big) \in \mathcal{S}(\mathbb{R}^{(e^{N+1}-1)/(e-1)}). \tag{7.4}$$

[2] And v will likely depend on the length of the input to ℓ.

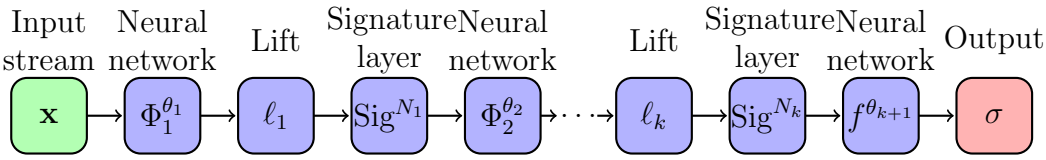

Figure 7.2: Deep signature model. Trainable parameters: $\theta_1, \ldots, \theta_{k+1}$.

In the above example,

$$\ell(\mathbf{x}) = (\mathbf{x}_2, \ldots, \mathbf{x}_n). \tag{7.5}$$

Other plausible choices are to cut up $\mathbf{x}$ into multiple pieces, for example

$$\ell(\mathbf{x}) = ((x_1, x_2), (x_3, x_4), \ldots, (x_{2\lfloor n/2 \rfloor - 1}, x_{2\lfloor n/2 \rfloor})), \tag{7.6}$$

or to take a sliding window

$$\ell(\mathbf{x}) = ((x_1, x_2, x_3), (x_2, x_3, x_4), \ldots, (x_{n-2}, x_{n-1}, x_n)). \tag{7.7}$$

Now we may repeat the whole process as many times as desired. That is, suppose we wish to learn a map from $\mathcal{S}(\mathbb{R}^d)$ to $\mathcal{X}$, where $\mathcal{X}$ is some set, which may be finite for a classification problem or infinite for a regression problem. Let $c_i, d_i, e_i, N_i \in \mathbb{N}$ be such that $d_1 = d$ and $d_{i+1} = (c_i^{N_i+1} - 1)/(c_i - 1)$, for $i = 1, \ldots, k$.

Let

$$\Phi_i^{\theta_i} \colon \mathbb{R}^{d_i \times m_i} \to \mathbb{R}^{e_i}, \qquad \ell_i \colon \mathcal{S}(\mathbb{R}^{e_i}) \to \mathcal{S}(\mathcal{S}(\mathbb{R}^{c_i})), \qquad f^{\theta_{k+1}} \colon \mathcal{S}(\mathbb{R}^{(c_k^{N_k+1}-1)/(c_k-1)}) \to \mathcal{X},$$

where $\Phi_i^{\theta_i}$ and ℓ_i are defined in the manner of equations (7.1)–(7.3) and (7.5)–(7.7), and $\theta_1, \ldots, \theta_{k+1}$ are some trainable parameters. Then defining compositions in the manner of equations (7.1)–(7.4), let

$$\sigma = \left(f^{\theta_{k+1}} \circ \mathrm{Sig}^{N_k} \circ \ell_k \circ \Phi_k^{\theta_k} \circ \cdots \circ \Phi_2^{\theta_2} \circ \mathrm{Sig}^{N_1} \circ \ell_1 \circ \Phi_1^{\theta_1} \right)(\mathbf{x}).$$

This defines the *deep signature model*, summarised in Figure 7.2. The name comes from the fact that the signature is now embedded deep within the network instead of

operating as a feature transformation. When there are multiple signature layers in a network, so that signatures-of-signatures are computed, then the name may be taken to have a double meaning, but we stress that this need not be the case. Good results may often be obtained with only a single signature layer.

An important special case is when $\mathcal{X} = \mathcal{S}(\mathbb{R}^e)$, so that the final network $f^{\theta_{k+1}}$ is stream-preserving. Then the overall model $\mathbf{x} \to \sigma$ is also stream-preserving. See for example Section 7.4.

In principle it is also fine to take the trivial lift of a sequence of a single element,

$$\ell(\mathbf{x}) = (\mathbf{x}). \tag{7.8}$$

Taking the signature of this will then essentially remove the stream-like nature, however, so it is suitable only for the final signature of a deep signature model. We observe in particular that this is what is done in the models described in Figure 7.1, which we identify as very special cases of the deep signature model, also lacking any learned transformation before the signature.

It is easy to see that the deep signature model exhibits the universal approximation property provided that the model is large enough.

7.3.4 Inverting the truncated signature

How well does a truncated signature encode the original stream of data? A simple experiment is to attempt to recover the original stream of data given its truncated signature. We remark that finding a mathematical description of this inversion is a challenging task [106, 30, 29].

Fix a stream of data $\mathbf{x} = (x_1, \ldots, x_n) \in \mathcal{S}(\mathbb{R}^d)$. Assume that the truncated signature $\mathrm{Sig}^N(\mathbf{x})$ and the number of steps $n \in \mathbb{N}$ are known. Now apply gradient descent to minimise

$$L(\mathbf{y}; \mathbf{x}) = \left\| \mathrm{Sig}^N(\mathbf{y}) - \mathrm{Sig}^N(\mathbf{x}) \right\|_2^2 \quad \text{for } \mathbf{y} = (y_1, \ldots, y_n) \in \mathcal{S}(\mathbb{R}^d).$$

Figure 7.3: Original path (blue) and path reconstructed from its signature (dashed orange) for four handwritten digits in the PenDigits dataset [45].

Figure 7.3 shows four handwritten digits from the PenDigits dataset [45]. The solid blue path is the original path $\mathbf{x}$, whilst the dashed orange path is the reconstructed path $\mathbf{y}$ minimising $L(\mathbf{y}; \mathbf{x})$. Truncated signatures of order $N = 12$ were used for this task. We see that the truncated signatures have managed to encode the input paths $\mathbf{x}$ almost perfectly.

7.4 A generative model for a stochastic process

Generative models are typically trained to learn to transform random noise to a target distribution. One common approach are Generative Adversarial Networks [63]. Here, the objective is to simultaneously train two models: a generative model and a discriminative model. The generative model's aim is to generate samples that look like they come from the same distribution as the data. It tries to *fool* the discriminative model, whose task is to determine whether a sample or set of samples come from the data or not.

An alternative approach is to define a distance on the space of distributions by embedding them into a Reproducing Kernel Hilbert Space. The discriminator is then a fixed two-sample test based on a kernel maximum mean discrepancy. This is known as a Generative Moment Matching Network [98, 49, 68].

With this framework we propose a deep signature model to generate sequential data. The discriminator is as in [90, 34]. The natural choice for random noise is Brownian motion B_t.

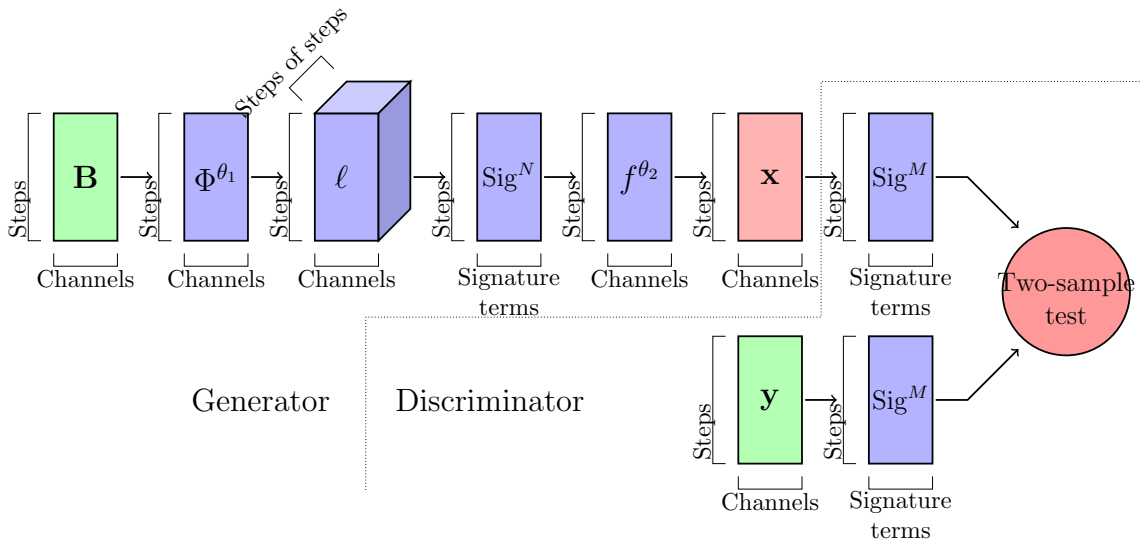

Figure 7.4: Generative model architecture. Trainable parameters: θ_1, θ_2. There is an implicit batch dimension throughout.

Define the kernel $k\colon \mathcal{S}(\mathbb{R}^d) \times \mathcal{S}(\mathbb{R}^d) \to \mathbb{R}$ by

$$k(\mathbf{x}, \mathbf{y}) = \left(\mathrm{Sig}^N\left(\lambda_{\mathbf{x}}\mathbf{x}\right), \mathrm{Sig}^N\left(\lambda_{\mathbf{y}}\mathbf{y}\right)\right),$$

where $\lambda_{\mathbf{x}} \in \mathbb{R}$ is a certain normalising constant which guarantees that k is the kernel of a Reproducing Kernel Hilbert Space, and $(\,\cdot\,,\,\cdot\,)$ denotes the dot product.

Given n samples $\{\mathbf{x}^{(i)}\}_{i=1}^n \subseteq \mathcal{S}(\mathbb{R}^d)$ from the generator and m samples $\{\mathbf{y}^{(i)}\}_{i=1}^m \subseteq \mathcal{S}(\mathbb{R}^d)$ from the target distribution, define the loss T by

$$T\left(\{\mathbf{x}^{(i)}\}_{i=1}^n, \{\mathbf{y}^{(i)}\}_{i=1}^m\right) = \frac{1}{n^2}\sum_{i,j} k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) - \frac{2}{nm}\sum_{i,j} k(\mathbf{x}^{(i)}, \mathbf{y}^{(j)}) + \frac{1}{m^2}\sum_{i,j} k(\mathbf{y}^{(i)}, \mathbf{y}^{(j)}). \tag{7.9}$$

Given two stream-preserving neural networks Φ^{θ_1} and f^{θ_2}, and a lift ℓ, then the generative model is defined by

$$\mathbf{x} = (f^{\theta_2} \circ \mathrm{Sig}^N \circ \ell \circ \Phi^{\theta_1})(\mathbf{B}),$$

where the input to the network is time-augmented Brownian motion

$$\mathbf{B} = ((t_1, B_{t_1}), \ldots, (t_n, B_{t_n})) \in \mathcal{S}(\mathbb{R}^2).$$

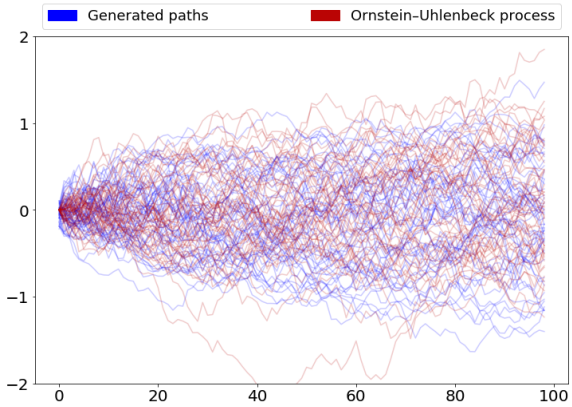

Figure 7.5: Generated paths alongside the original paths.

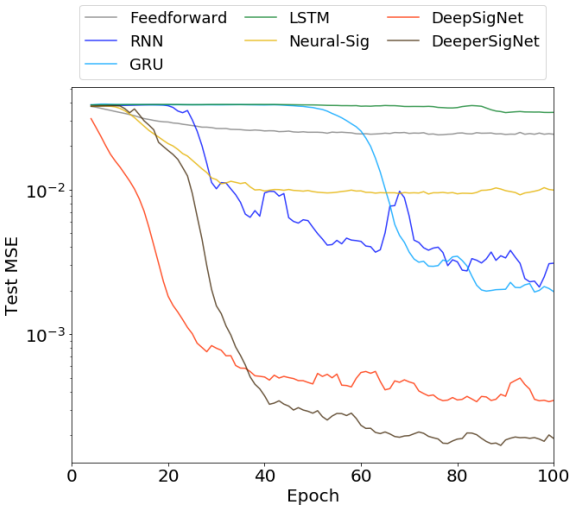

Figure 7.6: Out-of-sample performance at estimating the Hurst parameter for various models, with and without signatures, for a particular training run.

The overall model is shown in Figure 7.4. In a nice twist, both the generator and the discriminator involve the signature.

Observe how the generative part is a particular case of the deep signature model, and that furthermore the whole generator-discriminator pair is also a particular case of the deep signature model, with the trivial lift of equation (7.8) before the second signature layer.

We applied the proposed generative model to a dataset of 1024 realisations of an Ornstein–Uhlenbeck process [130]. The loss 7.9 was minimised at 6.6×10^{-4}. Applying the statistical test [34], this loss means that with a confidence level of 99% we cannot reject that the generated paths come from the real Ornstein–Uhlenbeck process. Figure 7.5 shows the generated paths alongside the original ones. Further implementation details are in Appendix E.

7.5 Supervised learning with fractional Brownian motion

Fractional Brownian motion [110] is a Gaussian process $B^H \colon [0, \infty) \to \mathbb{R}$ that generalises Brownian motion. It is self-similar and exhibits fractal-like behaviour. Fractional Brownian motion depends upon a parameter $H \in (0, 1)$, known as the Hurst parameter. Lower Hurst parameters result in noticeably rougher paths. The case of $H = 1/2$ corresponds to usual Brownian motion. Fractional Brownian motion has been successfully used to model phenomena in diverse fields. For example, empirical evidence from financial markets [60] suggests that log-volatility is well modelled by fractional Brownian motion with Hurst parameter $H \approx 0.1$.

Estimating the Hurst parameter of a fractional Brownian motion path is considered a nontrivial task because of the paths' non-stationarity and long-range dependencies [91]. We train a variety of models to perform this estimation. That is, to learn the map $\mathbf{x}^H \mapsto H$, where

$$\mathbf{x}^H = ((t_0, B^H_{t_0}), \ldots, (t_n, B^H_{t_n})) \in \mathcal{S}(\mathbb{R}^2)$$

for some realisation of B^H.

The results are shown in Figure 7.6 and Table 7.1. Also shown are the results of the rescaled range method [82], which is a mathematically derived method rather than a learned method. The data was flattened and fed into a non-recurrent feedforward model, referred to as Feedforward, as a baseline in the context of neural networks, whilst the simple Neural-Sig model outlined previously in Figure 7.1a provides a baseline from the context of signatures. DeepSigNet and DeeperSigNet are both deep signature models, of the form outlined by Figure 7.2. DeepSigNet has a single large Neural-Lift-Signature block, whilst DeeperSigNet has three smaller ones. We observe

that the deep signature models outperform all other models by at least an order of magnitude. Further implementation details are found in Appendix E.

	Test MSE		
	Mean	Variance	# Params
Rescaled Range	7.2×10^{-2}	3.7×10^{-3}	N/A
LSTM	4.3×10^{-2}	8.0×10^{-3}	12961
Feedforward	2.8×10^{-2}	3.0×10^{-3}	10209
Neural-Sig	1.1×10^{-2}	8.2×10^{-4}	10097
GRU	3.3×10^{-3}	1.3×10^{-3}	9729
RNN	1.7×10^{-3}	4.9×10^{-4}	10091
DeepSigNet	2.1×10^{-4}	8.7×10^{-5}	9261
DeeperSigNet	$\mathbf{1.6 \times 10^{-4}}$	2.1×10^{-5}	9686

Table 7.1: Final test mean squared error (MSE), averaged over 3 training runs, ordered from largest to smallest.

7.6 Non-Markovian deep reinforcement learning

Finally we show how these ideas may be extended by demonstrating a ResNet style adaptation of the deep signature model. This model uses signatures as the memory of a recurrent neural network. As an example, we demonstrate how this architecture may be used to tackle a non-Markovian reinforcement learning problem. Many tasks in reinforcement learning are Markovian, in the sense that the optimal action for the agent to take only depends on the current state of the environment, and not the past history of states. However if this is not the case then the task is usually harder to solve [134], and the agent is then required to maintain a memory in order to retain information about past states.

Let $\Phi^{\theta_1} \colon \mathbb{R}^d \to \mathbb{R}^e$ and $f^{\theta_2} \colon \mathbb{R}^{d+(e^{N+1}-1)/(e-1)} \to \{\text{actions}\}$ be functions depending on learnable parameters θ_1, θ_2. Given input $x_i \in \mathbb{R}^d$ at time i, let

$$y_i = \Phi^{\theta_1}(x_i), \qquad \sigma_i = \sigma_{i-1} \otimes \mathrm{Sig}^N((y_{i-1}, y_i)), \qquad a_i = f^{\theta_2}(x_i, \sigma_i),$$

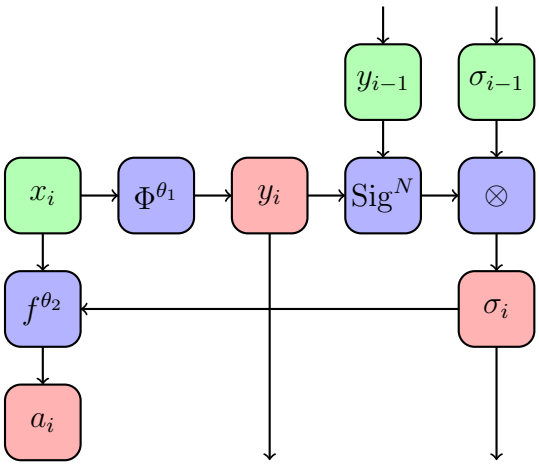

Figure 7.7: Agent architecture. Trainable parameters: θ_1, θ_2.

where $\otimes$ denotes the tensor product, a_i is the action proposed by the network at time i, and y_i and σ_i are the memory at time i. In particular this means that

$$\sigma_i = \mathrm{Sig}^N(\Phi^{\theta_1}(x_1), \dots, \Phi^{\theta_1}(x_i)) \in \mathbb{R}^{(e^{N+1}-1)/(e-1)}.$$

The model is summarised in Figure 7.7. Note that y_i is preserved in memory only to compute the signature at the next time step, as only paths of length at least two have meaningful signatures.

The inputs x_i may be collected into a stream $(x_i)_i \in \mathcal{S}(\mathbb{R}^d)$, the memory y_i may be collected into a stream $(y_i)_i \in \mathcal{S}(\mathbb{R}^e)$, and the memory σ_i may be collected into a stream $(\sigma_i)_i \in \mathcal{S}(\mathbb{R}^{(e^{N+1}-1)/(e-1)})$. Thus we see the relation to the general deep signature model. This model has a single Neural-Lift-Signature block, with a skip connection across the whole block. The neural component is given by the neural network Φ^{θ_1}, which is stream-preserving as it operates pointwise, in the manner of equation (7.1). The lift is the 'expanding window' lift given by equation (7.5). Finally f^{θ_2} is another neural network, which is again pointwise and thus stream-preserving.

We test this model on a non-Markovian modification to the classical Mountain Car problem [19], in which the agent receives only partial information: it is only given the

car's position, and not its velocity.[3] We find that it is capable of learning how to solve the problem within a set number of episodes, whilst a comparable RNN architecture fails to do so. The reinforcement learning technique used was Deep Q Learning [111] with the specified models performing function approximation on Q. Both models were chosen to have comparable numbers of parameters. Further implementation details can be found in Appendix E.

[3]Thus this belongs to particular class of non-Markov problems called Hidden Markov Models [118].

Chapter 8

A signature-based machine learning model for distinguishing bipolar disorder and borderline personality disorder

This chapter shows a real-world application of signatures. The chapter is based on the paper [117], published at *Translational Psychiatry*. In the paper, we studied whether signatures can be used to obtain information about individuals based on a sequence of observations about the person, in the context of psychiatry. Given that the signature provides a compact and effective summary of a path, we showed that they can be effectively used to make predictions about individuals, even with a rather small dataset. The paper was written in collaboration with Guy M. Goodwin, John R. Geddes and Kate E. A. Saunders from the Department of Psychiatry at the University of Oxford, and Terry Lyons. The code that was used in the paper can be found in `github.com/alan-turing-institute/signatures-psychiatry`.

8.1 Introduction

Historically the diagnosis of psychiatric disorders has been hampered by the inherent inaccuracy of retrospective narrative recall of abnormal mental states and the difficulty of defining their persistence over time. Furthermore, diagnostic classifications

have changed little since the original description made by Kraepelin [50] and are not clearly related to the underlying biology, which remains unknown ([76]). The shortcomings of current diagnostic categories have motivated attempts (notably the NIMH Research Domain Criteria ['RDoC']) to take a radical and 'bottom up' approach to diagnosis. In recent years the rapid development of mobile technologies has allowed more precise measures of subjective psychopathology. However analysis of these data streams has been challenging. Often the data generated by these measurements is sequential in nature, where the most valuable information is contained in the order at which different events occur, rather than in any individual event. These types of data need an analytic approach that exploits this distinctive feature. Signatures from rough path theory ([105]) have proven to be an effective way of analysing these streams of data since they naturally capture the order of events. As a consequence, signatures have successfully been applied to a number of machine learning problems where the object that is being studied is a stream of data. For instance, using the path signature as the core feature representation for handwritten sequential data, Dr. B. Graham won the ICDAR 2013 international challenge on online Chinese character recognition ([65]). The effectiveness of signatures in data science has lead to a number of publications where a signature-based approach has achieved better performance relative to approaches that do not use signatures ([65, 135, 96, 92, 141]).

Bipolar disorder (BD) and borderline personality disorder (BPD) are common mental disorders ([35, 66, 108]). BD is a mood disorder with a strong genetic basis while BPD is a disorder of personality commonly related to abusive experiences in childhood. Although the two disorders are thought to develop through different processes and mechanisms, BD and BPD can be difficult to differentiate in the clinic because both are characterised by mood instability and impulsivity and the behavioural consequences can be similar ([122]). However, even in the current state of knowledge,

correct diagnosis is essential because of the contrasting treatment approaches to the two disorders ([55, 54]). Here we use a signature-based machine learning model to analyse data obtained from a clinical study ([128]), which explored daily reporting of mood in participants with bipolar disorder, borderline personality disorder and healthy volunteers. Specifically, we sought to classify the diagnosis of participants on the basis of their evolving mood. The generality of the signature-based machine learning model allows these problems to be treated in similar ways.

8.2 Material and methods

8.2.1 Data

Prospective mood data was captured from 130 individuals who were taking part in the AMoSS study. The AMoSS study was a prospective longitudinal study that used a range of wearables in combination with a bespoke smartphone app to better characterise mood instability and its correlates in bipolar disorder ($N = 48$), borderline personality disorder ($N = 31$) and healthy volunteers ($N = 51$). Participants rated their mood daily across six different categories (anxiety, elation, sadness, anger, irritability and energy) using a 7-point Likert scale (with values from 1 (not at all) to 7 (very much)). Data were collected from each participant for a minimum of 3 months, although 61 of the 130 participants provided data for more than 12 months. Overall compliance was 81.2%. The mood data was divided into buckets comprised of streams of 20 consecutive observations. Although these observations were typically recorded daily, data did not have to come from 20 consecutive days. Thus, streams could be constructed even when participants failed to record their mood on some days. This generated 733 buckets of streamed data. These were randomly separated into a training set of 513, and a testing set of 220 streams of data. We then normalised the data to capture the intrinsic properties of the data independent of arbitrary units such as

the scale used to record the scores or the period of time when the study was carried out. The stream of mood reports was identified as a 7-dimensional path (1 dimension for time, 6 dimensions for the scores) defined on the unit interval. Figure 8.1 shows the aggregated normalised anxiety scores of a participant with bipolar disorder plotted against normalised time. Figure 8.2 shows the normalised scores of each category plotted against all other categories. This way of representing the path allows interpretation of the order at which the participant changes mood. Take for example, the Angry vs Elated plot (third row, first column). The path starts at the point (0, 0). As we see, the path moves left first. Therefore, the participant is becoming less and less elated, while the score in anger remains approximately constant. Suddenly, the period of low elation stops, and the participant starts recording low scores in anger. This low level in anger remains persistent for the rest of the path. On the other hand, if one considers the Angry vs Irritable plot (fifth row, first column), we observe that it is essentially a straight line and the levels of anger and irritability match closely for this participant.

8.2.2 Group classification

We started by establishing a set of input-output pairs $\{(R_i, Y_i)\}$, where R_i is the 7-dimensional path of participant i, as described above, and Y_i denotes the group it was diagnosed into at the beginning of the study. This set was transformed into a new set of input/output pairs, $\{(S^n(R_i), Y_i)\}_i$ where $S^n(R_i)$ denotes the truncated signature ([105]) of order n of the stream R_i. The truncated signature is a feature set describing the stream whose size depends on the truncation degree but not on the number of sample points (or indeed the paramaterisation) of the stream. The signature of order n grows exponentially with n, and therefore large values of n will produce input vectors of large dimensions but far smaller than other feature sets where the sampling rate

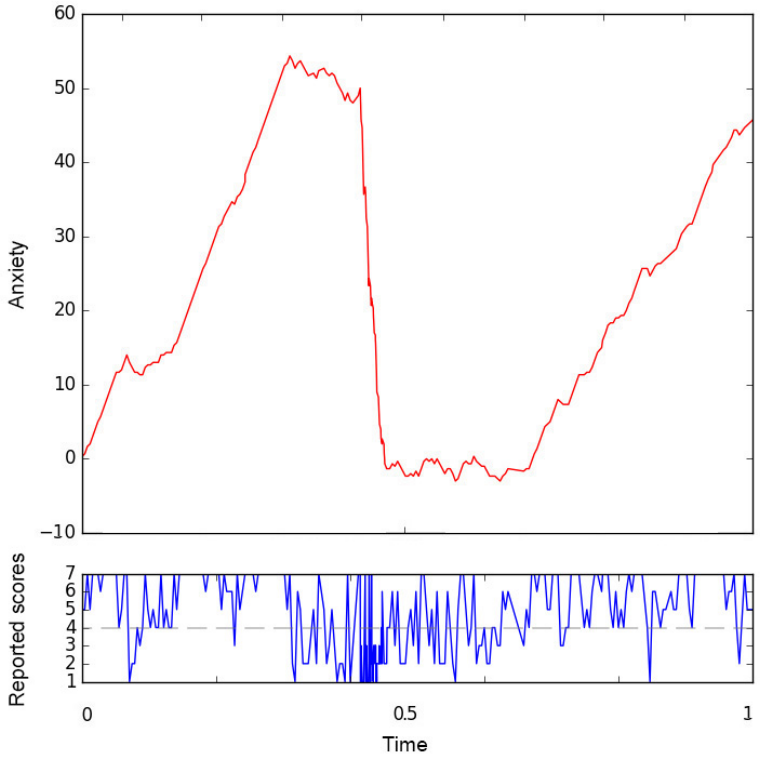

Figure 8.1: Normalised anxiety scores of a participant with bipolar disorder (above), which were calculated by aggregating the reported scores (below) plotted against normalised time. As we see, high levels of reported scores correspond to upward trends, low levels of reported scores correspond to downward trends and periods of time of high oscillations in the reported scores are represented by oscillations in the path.

directly changes the dimension of the feature set. We have only considered feature sets based on truncated signatures up to degree $n = 1, 2, 3, 4$. Random forest ([18]) was used to learn the mapping between the input and the output. Given that it has already been established that there are differences in mean mood scores between the groups overall ([128]) we also calculated the mean score in each mood category and classified streams of 20 consecutive observations on this basis as a comparison to the signature method. To assess the performance of the classification procedure we computed accuracy, sensitivity, specificity, positive predicted value (PPV). We also trained the model using pairs of clinical groups and computing the area under

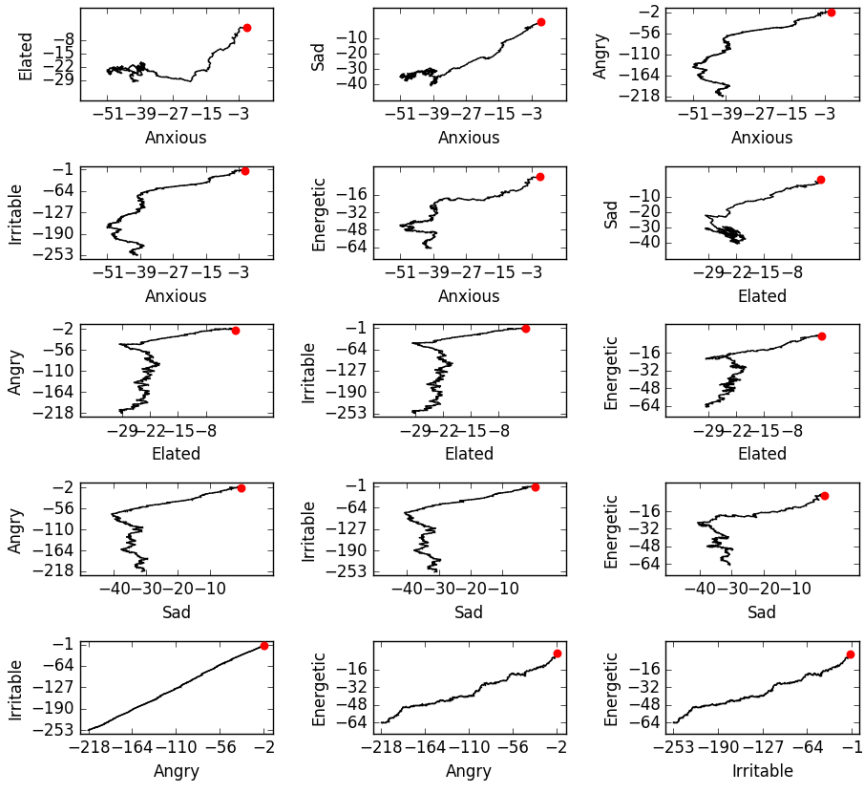

Figure 8.2: Normalised scores of each category plotted against all other categories, for a participant with bipolar disorder. The red point indicates the starting point. Notice that the scale is different in each plot.

the receiver operating characteristic curve (ROC) to assess the performance of the classification model at different values of threshold. Since this was done for each pair of clinical groups, we obtained a matrix with different values of the AUC. Moreover, in order to have a better understanding of how robust these percentages are, we applied bootstrapping to the model with the signature order fixed to 2. We also considered the extent to which participants were characterised as belonging to each specific clinical group. We trained the model using all participants except the one we were interested in, and then tested the model with 20-observations periods from this person. We calculated the proportion of periods of time when the participant was classified as bipolar, healthy or borderline, which allowed us to plot the participant as a point in the triangle. This process was followed for every participant, although we

excluded participants that generated less than 5 buckets of 20 observations in order to avoid extreme values for them. 121 participants satisfied this criterion.

8.3 Results

The signature-based model categorised 75% of participants into the correct diagnostic group when the second order signature was considered ($n = 2$, where n denotes the order of the signature). When the order was increased, the accuracy dropped slightly (70% for $n = 3$ and 69% for $n = 4$). The signature method performed significantly better than the naive model using the mean score in each category over the 20 observations, which classified just 54% of participants correctly. On bootstrapping the accuracy of the second order signature was 74.85% while the standard deviation was 2.05, suggesting that the results are stable and robust.

The signature of order 2 was able to capture second-order effects on the stream of data and since the feature set has 57 dimensions it was reasonable to learn linear relations from the 517 buckets of data in the training set. To capture the full second-order effects using raw streams of data, the feature set would have required over 10,000 dimensions. A 57-dimensional feature set of randomly selected quadratic polynomials on the raw data stream was inferior to the signature model: the average accuracy was 68% for the random feature set, well below the 75% accuracy of the second- order signature model.

The proportion of a participant's sequential mood buckets matched each specific diagnostic group and positioned that specific individual on a triangular spectrum; the distribution of that density for each initial diagnosis is shown in Figure 8.3. In each of the plots, the regions of highest density of participants are located in the correct corner of the triangle: the greatest consistency is with the healthy participants.

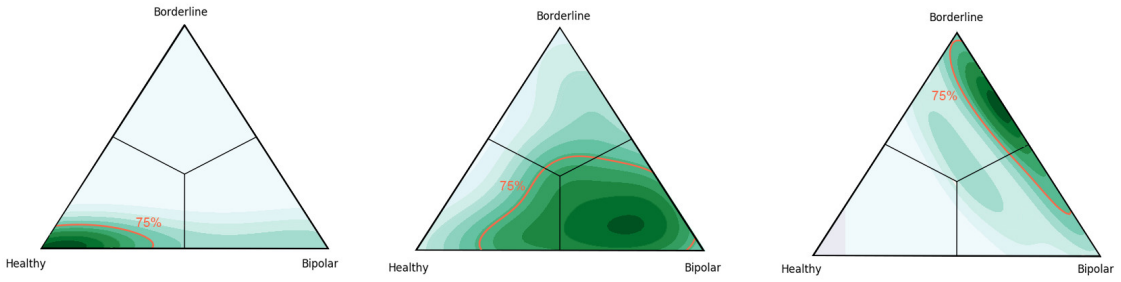

Figure 8.3: From left to right: healthy participants, bipolar participants and borderline participants.

	Healthy		Bipolar		Borderline	
	Accuracy	AUC	Accuracy	AUC	Accuracy	AUC
Healthy	-	-	84%	0.91	93%	0.98
Bipolar	-	-	-	-	80%	86%
Borderline	-	-	-	-	-	-

Table 8.1: Accuracy and are under the ROC curve.

Table 8.1 shows that the signature-based model clearly distinguishes healthy participants from the clinical groups, obtaining an accuracy of 84% in classifying BD participants and HC, and 93% in classifying BPD participants and HC.

8.4 Discussion

This analysis uses the signature feature set combined with random forests to study self-reported mood data. The signature-based model, with its ability to ignore parameterisation and provide canonical low dimensional sets of features, proved to be effective in distinguishing the three participant groups on the basis of their self-reported mood and was more accurate in doing so than previously used metrics, specifically mean mood scores. Furthermore, the method showed a clear overlap between the HC and BD groups suggestive of a normalization of mood structure in some BD participants for some of the time; this confirms clinical experience.

By contrast there is a very clear differentiation between HC and BPD. This finding

supports the concept of BPD as a distinct disorder. The distinctness from the healthy controls makes it more difficult to conceptualize BPD as an extreme variation of normal personality. There was some overlap between BD and BPD, but also marked differences. Overall the signatures of mood in the BPD group do appear to be distinct from both health and bipolar groups. The findings have implications for diagnostic practice where rates of misdiagnosis of BPD and BD are significant ([122, 142]). There is previous work published on machine learning approaches to distinguishing BD and BPD from healthy controls. For instance, some publications ([122, 142]) report that MRI data can be used to classify individuals with borderline personality disorder and healthy controls, obtaining an accuracy of and 80% ([123]) and 93.55% ([136]), although the latter percentage should be taken with caution due to the small sample size. Neuroimaging has also been used to detect bipolar disorder ([124]) with an accuracy of 59%, as well as other sources of data such as verbal fluency ([37]) with an accuracy of 79%. Therefore, our approach has improved existing results, although comparison across the different papers is difficult due to the different datasets that were used.

Our results suggest that the second order signature is the most effective one given the available data. The signature of second order is able to capture second order effects across pairs of moods in a meaningful and concise way. If the same effect wanted to be obtained using the raw parameterised 7-dimensional path, we would have an initial feature of more than 10,000 dimensions – which and to train usefully would have required a much larger data set than our 733 instances. There is a clear trade-off between choosing a high signature order, which corresponds to a richer model but may result in overfitting depending on the size of the dataset, and choosing a small signature order, which avoids overfitting but results in a simpler model.

8.4.1 Limitations

The study population was derived from patient samples but did not include those who were acutely unwell or who had significant comorbidities and as such may have represented a more stable sub-population of patients. We had relatively little contact with patients during the study and only a baseline clinical assessment was carried out such that the validity of the self-reported mood scores cannot be certain. However, we have no evidence to suggest that the reports are inaccurate and they are likely to be a good representation of the patients' lived experience. This is supported by the fact that, as shown in this paper, the self-reported scores were significant enough to meaningfully distinguish the three groups. We also kept no ongoing record of medication changes through the duration of the study.

Our dataset consisted of 733 examples. The dataset was quite small for machine learning purposes and so we avoided testing different models in order to avoid over-fitting. A bigger dataset would undoubtedly allow more flexibility.

Bibliography

[1] Beatrice Acciaio, Mathias Beiglböck, Friedrich Penkner, and Walter Schachermayer. A model-free version of the fundamental theorem of asset pricing and the super-replication theorem. *Mathematical Finance*, 26(2):233–251, 2016.

[2] Y. Aït-Sahalia and A. W. Lo. Nonparametric estimation of state-price densities implicit in financial asset prices. *The Journal of Finance*, 53(2):499–547, 1998.

[3] Robert Almgren and Neil Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3:5–40, 2001.

[4] A. Ananova and R. Cont. Pathwise integration with respect to paths of finite quadratic variation. *Journal de Mathématiques Pures et Appliquées*, 107(6):737–757, 2017.

[5] K. J. Arrow. The role of securities in the optimal allocation of risk-bearing. In *Readings in Welfare Economics*, pages 258–263. Palgrave, London, 1973.

[6] Emmanuel Bacry, Adrian Iuga, Matthieu Lasnier, and Charles-Albert Lehalle. Market impacts and the life cycle of investors orders. *Market Microstructure and Liquidity*, 1(02):1550009, 2015.

[7] Vlad Bally, Lucia Caramellino, Rama Cont, Frederic Utzet, and Josep Vives. *Stochastic integration by parts and functional Itô calculus*. Springer, 2016.

[8] Ole E Barndorff-Nielsen and Neil Shephard. Modelling by Lévy processess for financial econometrics. In *Lévy processes*, pages 283–318. Springer, 2001.

[9] Ole E Barndorff-Nielsen and Neil Shephard. Non-Gaussian Ornstein–Uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(2):167–241, 2001.

[10] Kyle Bechler and Michael Ludkovski. Optimal Execution with Dynamic Order Flow Imbalance. *SIAM Journal on Financial Mathematics*, 6(1):1123–1151, 2015.

[11] Mathias Beiglböck, Alexander MG Cox, Martin Huesmann, Nicolas Perkowski, and David J Prömel. Pathwise superreplication via Vovk's outer measure. *Finance and Stochastics*, 21(4):1141–1166, 2017.

[12] Dimitris Bertsimas and Andrew W Lo. Optimal control of execution costs. *Journal of Financial Markets*, 1(1):1–50, 1998.

[13] Hans Bühler, Lukas Gonon, Josef Teichmann, and Ben Wood. Deep Hedging. *arXiv: Computational Finance*, 2018.

[14] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of political economy, 81*, 3:637–654, 1973.

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

[16] Bruno Bouchard, Marcel Nutz, et al. Arbitrage and duality in nondominated discrete-time models. *The Annals of Applied Probability*, 25(2):823–859, 2015.

[17] D. T. Breeden and R. H. Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of business*, pages 621–651, 1978.

[18] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.

[19] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI gym. *arXiv preprint arXiv:1606.01540*, 2016.

[20] Matteo Burzoni, Marco Frittelli, Zhaoxu Hou, Marco Maggis, and Jan Obloj. Pointwise arbitrage pricing theory in discrete time. *Mathematics of Operations Research*, 2019.

[21] Álvaro Cartea, Ryan Francis Donnelly, and Sebastian Jaimungal. Algorithmic Trading with Model Uncertainty. *Siam Journal on Financial Mathematics*, 8(1):635–671, 2017.

[22] Álvaro Cartea and Sebastian Jaimungal. Optimal execution with limit and market orders. *Quantitative Finance*, 15(8):1279–1291, 2015.

[23] Álvaro Cartea and Sebastian Jaimungal. A closed-form execution strategy to target volume weighted average price. *SIAM Journal on Financial Mathematics*, 7(1):760–785, 2016.

[24] Álvaro Cartea and Sebastian Jaimungal. Incorporating order-flow into optimal execution. *Mathematics and Financial Economics*, 10(3):339–364, 2016.

[25] Álvaro Cartea, Sebastian Jaimungal, and Tianyi Jia. Trading Foreign Exchange Triplets. *Available at SSRN 3054656*, 2017.

[26] Alvaro Cartea, Sebastian Jaimungal, and José Penalva. *Algorithmic and High-Frequency Trading*. 2015.

[27] Alvaro Cartea, Imanol Perez Arribas, and Leandro Sanchez-Betancourt. Optimal Execution of Foreign Securities: A Double-Execution Problem with Signatures and Machine Learning. *Available at SSRN: http://ssrn.com/abstract=3562251*.

[28] Philippe Casgrain, Brian Ning, and Sebastian Jaimungal. Deep Q-learning for Nash equilibria: Nash-DQN. *arXiv preprint arXiv:1904.10554*, 2019.

[29] Jiawei Chang. *Effective algorithms for inverting the signature of a path*. PhD thesis, University of Oxford, 2018.

[30] Jiawei Chang, Nick Duffield, Hao Ni, Weijun Xu, et al. Signature inversion for monotone paths. *Electronic Communications in Probability*, 22, 2017.

[31] Pierre Chen, Edmond Lezmi, Thierry Roncalli, and Jiali Xu. A note on portfolio optimization with quadratic transaction costs. *arXiv preprint arXiv:2001.01612*, 2020.

[32] I. Chevyrev and A. Kormilitzin. A primer on the signature method in machine learning. *arXiv preprint arXiv:1603.03788*, 2016.

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

[34] Ilya Chevyrev and Harald Oberhauser. Signature moments to characterize laws of stochastic processes. *arXiv preprint arXiv:1810.10971*, 2018.

[35] Jeremy Coid, Min Yang, Peter Tyrer, Amanda Roberts, and Simone Ullrich. Prevalence and correlates of personality disorder in Great Britain. *The British Journal of Psychiatry*, 188(5):423–431, 2006.

[36] R. Cont and D. A. Fournié. Functional Itô calculus and stochastic integral representation of martingales. *The Annals of Probability*, 41(1):109–133, 2013.

[37] Sergi G Costafreda, Cynthia HY Fu, Marco Picchioni, Timothea Toulopoulou, Colm McDonald, Eugenia Kravariti, Muriel Walshe, Diana Prata, Robin M Murray, and Philip K McGuire. Pattern of neural responses to verbal fluency shows diagnostic specificity for schizophrenia and bipolar disorder. *BMC psychiatry*, 11(1):18, 2011.

[38] Laure Coutin and Antoine Lejay. Semi-martingales and rough paths theory. *Electronic Journal of Probability*, 10(23):761–785, 2005.

[39] Laure Coutin and Zhongmin Qian. Stochastic analysis, rough path analysis and fractional Brownian motions. *Probability Theory and Related Fields*, 122(1):108–140, 2002.

[40] Gianbiagio Curato, Jim Gatheral, and Fabrizio Lillo. Optimal execution with non-linear transient market impact. *Quantitative Finance*, 17(1):41–54, 2017.

[41] Ngoc-Minh Dang. Optimal Execution with Transient Impact. *Market Microstructure and Liquidity*, 3(1):1750008, 2017.

[42] G. Debreu. *Theory of value: An axiomatic analysis of economic equilibrium.* Yale University Press, No. 17), 1987.

[43] F. Delbaen, P. Grandits, T. Rheinl"ander, D. Samperi, M. Schweizer, and C. Stricker. Exponential hedging and entropic penalties. *Mathematical finance, 12*, 2:99–123, 2002.

[44] Giovanni B Di Masi, Yu M Kabanov, and Wolfgang J Runggaldier. Mean-variance hedging of options on stocks with Markov volatilities. *Theory of Probability & Its Applications*, 39(1):172–182, 1995.

[45] Dheeru Dua and Casey Graff. UCI Machine Learning Repository, 2017.

[46] D. Duffie and H. R. Richardson. Mean-variance hedging in continuous time. *The Annals of Applied Probability*, 1(1):1–15, 1991.

[47] B. Dumas, J. Fleming, and R. E. Whaley. Implied volatility functions: Empirical tests. *The Journal of Finance*, 53(6):2059–2106, 1998.

[48] B. Dupire. Functional Itô calculus, 2009.

[49] Gintare Karolina Dziugaite, Daniel M. Roy, and Zoubin Ghahramani. Training generative neural networks via maximum mean discrepancy optimization. *UAI*, 2015.

[50] Andreas Ebert and Karl-Jurgen Bar. Emil kraepelin: A pioneer of scientific understanding of psychiatry and psychopharmacology. *Indian journal of psychiatry*, 52(2):191, 2010.

[51] Paul Embrechts. *Selfsimilar processes*, volume 21. Princeton University Press, 2009.

[52] G. Flint, B. Hambly, and T. Lyons. Discretely sampled signals and the rough Hoff process. *Stochastic Processes and their Applications*, 126(9):2593–2614, 2016.

[53] H. S. Foellmer. Hedging of non-redundant contingent claims. *(No*, 3, 1985.

[54] National Collaborating Centre for Mental Health et al. Bipolar disorder: the assessment and management of bipolar disorder in adults, children and young people in primary and secondary care. *Commissioned by the National Institute for Health and Care Excellence. Leicester and London: The British Psychological Society and the Royal College of Psychiatrists*, 2014.

[55] National Collaborating Centre for Mental Health (UK et al. Borderline personality disorder: treatment and management. British Psychological Society, 2009.

[56] P. K. Friz and N. B. Victoir. *Multidimensional stochastic processes as rough paths: theory and applications*, volume 120). Cambridge University Press, 2010.

[57] Jean-François Le Gall. A Path-Valued Markov Process and its Connections with Partial Differential Equations. pages 185–212, 1994.

[58] Jean-François Le Gall. A Path-Valued Markov Process and its Connections with Partial Differential Equations. pages 185–212, 1994.

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

[60] Jim Gatheral, Thibault Jaisson, and Mathieu Rosenbaum. Volatility is Rough. *Quantitative Finance, 18:6, 933-949*, 2018.

[61] Jim Gatheral, Alexander Schied, and Alla Slynko. Transient Linear Price Impact and Fredholm Integral Equations. *Mathematical Finance*, 22(3):445–474, 2012.

[62] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

[63] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

[64] S. Goutte, N. Oudjane, and F. Russo. Variance optimal hedging for continuous time additive processes and applications. *Stochastics An International Journal of Probability and Stochastic Processes, 86*, 1:147–185, 2014.

[65] B. Graham. Sparse arrays of signatures for online character recognition. arxiv. preprint, 2013.

[66] Bridget F Grant, S Patricia Chou, Risë B Goldstein, Boji Huang, Frederick S Stinson, Tulshi D Saha, Sharon M Smith, Deborah A Dawson, Attila J Pulay, Roger P Pickering, et al. Prevalence, correlates, disability, and comorbidity of DSM-IV borderline personality disorder: results from the Wave 2 National Epidemiologic Survey on Alcohol and Related Conditions. *The Journal of clinical psychiatry*, 69(4):533, 2008.

[67] M. R. Grasselli and T. R. Hurd. A Monte Carlo method for exponential hedging of contingent claims. *arXiv preprint arXiv:math/0211383*, 2002.

[68] A. Gretton, K. M. Borgwardt, M. Rasch, B. Scholkopf, and A. J Smola. A kernel method for the two-sample problem. *Advances in Neural Information Processing Systems*, 2007.

[69] Olivier Guéant. Optimal Execution and Block Trade Pricing: A General Framework. *Applied Mathematical Finance*, 22(4):336–365, 2015.

[70] Olivier Guéant. *The financial mathematics of market liquidity: From optimal execution to market making*, volume 33. CRC Press, 2016.

[71] Olivier Guéant. Optimal market making. *Applied Mathematical Finance*, 24(2):112–154, 2017.

[72] Olivier Guéant, Charles-Albert Lehalle, and Joaquin Fernandez-Tapia. Optimal Portfolio Liquidation with Limit Orders. *SIAM Journal on Financial Mathematics*, 3(1):740–764, 2012.

[73] Olivier Guéant and Iuliia Manziuk. Deep Reinforcement Learning for Market Making in Corporate Bonds: Beating the Curse of Dimensionality. *Applied Mathematical Finance*, 26(5):387–452, 2019.

[74] L. G. Gyurkó, T. Lyons, M. Kontkowski, and J. Field. Extracting information from the signature of a financial data stream. *arXiv preprint arXiv:1307.7244*, 2014.

[75] Ben Hambly and Terry Lyons. Uniqueness for the signature of a path of bounded variation and the reduced path group. *arXiv preprint math/0507536*, 2005.

[76] Paul J Harrison, John R Geddes, and Elizabeth M Tunbridge. The emerging neurobiology of bipolar disorder. *Trends in neurosciences*, 2017.

[77] Steven L Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The review of financial studies*, 6(2):327–343, 1993.

[78] Ben Hoff. *The Brownian frame process as a rough path*. PhD thesis, University of Oxford, 2005.

[79] Zhaoxu Hou and Jan Obloj. Robust pricing–hedging dualities in continuous time. *Finance and Stochastics*, 22(3):511–567, 2018.

[80] F. Hubalek, J. Kallsen, and L. Krawczyk. Variance-optimal hedging for processes with stationary independent increments. *The Annals of Applied Probability, 16*, 2:853–885, 2006.

[81] John Hull and Alan White. The pricing of options on assets with stochastic volatilities. *The journal of finance*, 42(2):281–300, 1987.

[82] H.E. Hurst. The Long-Term Storage Capacity of Reservoirs. *Transactions of the American Society of Civil Engineers*, 1951.

[83] J. M. Hutchinson, A. W. Lo, and T. Poggio. A nonparametric approach to pricing and hedging derivative securities via learning networks. *The Journal of Finance, 49*, 3:851–889, 1994.

[84] James M Hutchinson, Andrew W Lo, and Tomaso Poggio. A nonparametric approach to pricing and hedging derivative securities via learning networks. *The Journal of Finance*, 49(3):851–889, 1994.

[85] M. Jeanblanc, M. Mania, M. Santacroce, and M. Schweizer. Mean-variance hedging via stochastic control and BSDEs for general semimartingales. *The Annals of Applied Probability*, 22(6):2388–2428, 2012.

[86] Jasdeep Kalsi, Terry Lyons, and Imanol Perez Arribas. Optimal execution with rough path signatures. *SIAM Journal on Financial Mathematics*, 2020.

[87] Patrick Kidger, Patric Bonnier, Imanol Perez Arribas, Cristopher Salvi, and Terry Lyons. Deep Signature Transforms. In *Advances in Neural Information Processing Systems*, pages 3099–3109, 2019.

[88] Patrick Kidger and Terry Lyons. Signatory: differentiable computations of the signature and logsignature transforms, on both CPU and GPU. *arXiv preprint arXiv:2001.00706*, 2020.

[89] D Kingma and J Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.

[90] Franz J Király and Harald Oberhauser. Kernels for sequentially ordered data. *Journal of Machine Learning Research*, 2019.

[91] Lucas Lacasa, Bartolo Luque, Jordi Luque, and Juan Carlos Nuno. The visibility graph: A new method for estimating the Hurst exponent of fractional Brownian motion. *EPL (Europhysics Letters)*, 86(3):30001, 2009.

[92] Songxuan Lai, Lianwen Jin, and Weixin Yang. Online signature verification using recurrent neural network and length-normalized path signature descriptor. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, volume 1, pages 400–405. IEEE, 2017.

[93] Charles-Albert Lehalle and Eyal Neuman. Incorporating Signals into Optimal Trading. *Finance and Stochastics*, 2019.

[94] Tim Leung and Xin Li. Optimal mean reversion trading with transaction costs and stop-loss exit. *International Journal of Theoretical and Applied Finance*, 18(03):1550020, 2015.

[95] Daniel Levin, Terry Lyons, and Hao Ni. Learning from the past, predicting the statistics for the future, learning an evolving system. *arXiv preprint arXiv:1309.0260*, 2013.

[96] Chenyang Li, Xin Zhang, and Lianwen Jin. LPSNet: a novel log path signature feature based hand gesture recognition framework. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 631–639, 2017.

[97] Chenyang Li, Xin Zhang, and Lianwen Jin. LPSNet: a novel log path signature feature based hand gesture recognition framework. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 631–639, 2017.

[98] Yujia Li, Kevin Swersky, and Richard Zemel. Generative moment matching networks. *ICML*, 2015.

[99] T. Lyons, H. Ni, and H Oberhauser. A feature set for streams and an application to high-frequency financial tick data. *ICBDC*, 2014.

[100] Terry Lyons. Rough paths, Signatures and the modelling of functions on streams. *arXiv preprint arXiv:1405.4537*, 2014.

[101] Terry Lyons, Sina Nejad, and Imanol Perez Arribas. Nonparametric pricing and hedging of exotic derivatives. *arXiv preprint arXiv:1905.00711*, 2019.

[102] Terry Lyons, Sina Nejad, and Imanol Perez Arribas. Numerical method for model-free pricing of exotic derivatives using rough path signatures. *Applied Mathematical Finance*, 2019.

[103] Terry Lyons, Hao Ni, et al. Expected signature of Brownian motion up to the first exit time from a bounded domain. *The Annals of Probability*, 43(5):2729–2762, 2015.

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

[105] Terry J Lyons, Michael Caruana, and Thierry Lévy. *Differential equations driven by rough paths*. Springer, 2007.

[106] Terry J Lyons and Weijun Xu. Inverting the signature of a path. *Journal of the European Mathematical Society*, 20(7):1655–1687, 2018.

[107] Iacopo Mastromatteo, Michael Benzaquen, Zoltan Eisler, and Jean-Philippe Bouchaud. Trading Lightly: Cross-Impact and Optimal Portfolio Execution. *arXiv preprint arXiv:1702.03838*, 2017.

[108] Kathleen R Merikangas, Hagop S Akiskal, Jules Angst, Paul E Greenberg, Robert MA Hirschfeld, Maria Petukhova, and Ronald C Kessler. Lifetime and 12-month prevalence of bipolar spectrum disorder in the National Comorbidity Survey replication. *Archives of general psychiatry*, 64(5):543–552, 2007.

[109] R. C. Merton. Theory of rational option pricing. *The Bell Journal of economics and management science*, pages 141–183, 1973.

[110] Yuliya Mishura. *Stochastic calculus for fractional Brownian motion and related processes*, volume 1929. Springer Science & Business Media, 2008.

[111] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.

[112] H. Ni. *The expected signature of a stochastic process*. Doctoral dissertation, University of Oxford, 2012.

[113] Brian Ning, Franco Ho Ting Ling, and Sebastian Jaimungal. Double deep Q-learning for optimal execution. *arXiv preprint arXiv:1812.06600*, 2018.

[114] Paolo Pasquariello and Clara Vega. Strategic Cross-Trading in the U.S. Stock Market. *Review of Finance*, 19(1):229–282, 2015.

[115] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. 2017.

[116] Imanol Perez Arribas. Derivatives pricing using signature payoffs. *arXiv preprint arXiv:1809.09466*, 2018.

[117] 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):274, 2018.

[118] Lawrence R Rabiner and Biing-Hwang Juang. An introduction to hidden Markov models. *ieee assp magazine*, 3(1):4–16, 1986.

[119] Jeremy Reizenstein and Benjamin Graham. The iisignature library: efficient calculation of iterated-integral signatures and log signatures. *arXiv preprint arXiv:1802.08252*, 2018.

[120] Candia Riga. A pathwise approach to continuous-time trading. *arXiv preprint arXiv:1602.04946*, 2016.

[121] Candia Riga. A pathwise approach to continuous-time trading. *arXiv preprint arXiv:1602.04946*, 2016.

[122] Camilo J Ruggero, Mark Zimmerman, Iwona Chelminski, and Diane Young. Borderline personality disorder and the misdiagnosis of bipolar disorder. *Journal of psychiatric research*, 44(6):405–408, 2010.

[123] João Ricardo Sato, Gerardo Maria de Araujo Filho, Thabata Bueno de Araujo, Rodrigo Affonsecca Bressan, Pedro Paulo de Oliveira, and Andrea Parolin Jackowski. Can neuroimaging be used as a support to diagnosis of borderline personality disorder? An approach based on computational neuroanatomy and machine learning. *Journal of psychiatric research*, 46(9):1126–1132, 2012.

[124] Hugo G Schnack, Mireille Nieuwenhuis, Neeltje EM van Haren, Lucija Abramovic, Thomas W Scheewe, Rachel M Brouwer, Hilleke E Hulshoff Pol, and René S Kahn. Can structural MRI aid in clinical classification? A machine learning study in two independent samples of patients with schizophrenia, bipolar disorder and healthy subjects. *Neuroimage*, 84:299–306, 2014.

[125] M. Schweizer. Mean-variance hedging for general claims. *The annals of applied probability*, pages 171–179, 1992.

[126] M. Schweizer. *Mean–Variance Hedging.* Encyclopedia of Quantitative Finance, 2010.

[127] Bence Toth, Yves Lemperiere, Cyril Deremble, Joachim de Lataillade, Julien Kockelkoren, and Jean-Philippe Bouchaud. Anomalous Price Impact and the Critical Nature of Liquidity in Financial Markets. *Physical Review X*, 1(2), 2011.

[128] A Tsanas, KEA Saunders, AC Bilderbeck, N Palmius, M Osipov, GD Clifford, GM Goodwin, and M De Vos. Daily longitudinal self-monitoring of mood variability in bipolar disorder and borderline personality disorder. *Journal of affective disorders*, 205:225–233, 2016.

[129] Gerry Tsoukalas, Jiang Wang, and Kay Giesecke. Dynamic Portfolio Execution. *Management Science*, 2017.

[130] George E Uhlenbeck and Leonard S Ornstein. On the theory of the Brownian motion. *Physical review*, 36(5):823, 1930.

[131] Christopher John Cornish Hellaby Watkins. Learning from delayed rewards. 1989.

[132] A Elizabeth Whalley. Optimal partial hedging of options with small transaction costs. *Journal of Futures Markets*, 31(9):855–897, 2011.

[133] A Elizabeth Whalley and Paul Wilmott. An asymptotic analysis of an optimal hedging model for option pricing with transaction costs. *Mathematical Finance*, 7(3):307–324, 1997.

[134] Steven D Whitehead and Long-Ji Lin. Reinforcement learning of non-Markov decision processes. *Artificial Intelligence*, 73(1-2):271–306, 1995.

[135] Z. Xie, Z. Sun, L. Jin, H. Ni, and T. Lyons. Learning spatial-semantic context with fully convolutional recurrent network for online handwritten Chinese text recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(8):1903–1917, 2018.

[136] Tingting Xu, Kathryn R Cullen, Alaa Houri, Kelvin O Lim, S Charles Schulz, and Keshab K Parhi. Classification of borderline personality disorder based on spectral power of resting-state fMRI. In *Conference proceedings:... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference*, volume 2014, page 5036. NIH Public Access, 2014.

[137] Weixin Yang, Lianwen Jin, and Manfei Liu. Chinese character-level writer identification using path signature feature, DropStroke and deep CNN. In

2015 13th International Conference on Document Analysis and Recognition (ICDAR), pages 546–550. IEEE, 2015.

[138] Weixin Yang, Lianwen Jin, and Manfei Liu. Deepwriterid: An end-to-end online text-independent writer identification system. *IEEE Intelligent Systems*, 31(2):45–53, 2016.

[139] Weixin Yang, Lianwen Jin, Hao Ni, and Terry Lyons. Rotation-free online handwritten character recognition using dyadic path signature features, hanging normalization, and deep neural network. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 4083–4088. IEEE, 2016.

[140] Weixin Yang, Lianwen Jin, Dacheng Tao, Zecheng Xie, and Ziyong Feng. DropSample: A new training method to enhance deep convolutional neural networks for large-scale unconstrained handwritten Chinese character recognition. *Pattern Recognition*, 58:190–203, 2016.

[141] Weixin Yang, Terry Lyons, Hao Ni, Cordelia Schmid, Lianwen Jin, and Jiawei Chang. Leveraging the Path Signature for Skeleton-based Human Action Recognition. *arXiv preprint arXiv:1707.03993*, 2017.

[142] Mark Zimmerman, Camilo J Ruggero, Iwona Chelminski, and Diane Young. Psychiatric diagnoses in patients previously overdiagnosed with bipolar disorder. *The Journal of clinical psychiatry*, 2010.

Appendix A

Proofs of Chapter 3

Lemma A.0.1 (Signature of a perturbed rough path)**.** *Let $\mathbb{X} \in G\Omega_p([0,T],\mathbb{R}^d)$ be a p-rough path with $p \in [2,3)$, and let $\varphi : C([0,T];\mathfrak{so}(d))$ be of bounded variation. Define the second-level perturbation $\mathbb{Y} := \mathbb{X} + \varphi \in G\Omega_p([0,T];\mathbb{R}^d)$. Define $z^1 := X, z^2 := \varphi$. Given $\mathbf{I} = \mathbf{i_1} \dots \mathbf{i_k} \in \{\mathbf{1},\mathbf{2}\}^k$, let*

$$a^{\mathbf{I}}_{s,t} := \int_{s \le u_1 \le \dots \le u_k \le t} dz^{\mathbf{i_1}}_{u_1} \otimes \dots \otimes dz^{\mathbf{i_k}}_{u_k}.$$

Then, for each $N \ge 1$ the level N signature of $\mathbb{Y}$ is given by

$$\mathbb{Y}^N = \sum_k \sum_{\substack{\mathbf{I}=\mathbf{i_1}\dots\mathbf{i_k} \in \{\mathbf{1},\mathbf{2}\}^k \\ i_1+\dots+i_k=N}} a^{\mathbf{I}}.$$

Proof. If $N = 1$, the sum above is reduced to $a^{(1)} = \int_0^T dX_{0,t} = X_{0,T} = Y_{0,T}$, so that the claim holds. Assume that the statement is true for $N-1$. We will show that it also holds for $N \ge 2$.

By [100, Lemma 4.6], $\mathbb{Y}^{\le N}$ satisfies the rough differential equation

$$d\mathbb{Y}^{\le N} = \sum_{i=1}^{d} \mathbb{Y}^{\le N} \otimes e_i dY^i_t, \quad \mathbb{Y}^{\le N}_0 = 1 \in T^{(N)}(\mathbb{R}^d).$$

By [56, Theorem 12.16], the solution of the above rough differential equation is also the solution of the rough differential equation with drift

$$d\mathbb{Y}^{\leq N} = \sum_{i=1}^{d} \mathbb{Y}^{\leq N} \otimes e_i dX^i + \sum_{1\leq i<j\leq d} \mathbb{Y}^{\leq N} \otimes [e_i, e_j] d\varphi^{i,j}, \quad \mathbb{Y}_0^{\leq N} = 1 \in T^{(N)}(\mathbb{R}^d)$$

where $[e_i, e_j] := e_i \otimes e_j - e_j \otimes e_i$ denotes the Lie bracket. Hence, the level N projection of $\mathbb{Y}^{\leq N}$ will satisfy

$$\begin{aligned}\mathbb{Y}_{s,t}^N &= \int_s^t \mathbb{Y}_{s,u}^{N-1} \otimes dX_u + \sum_{1\leq i<j\leq d} \int_s^t \mathbb{Y}_{s,u}^{N-2} \otimes [e_i, e_j] d\varphi_u^{i,j} \\ &= \int_s^t \mathbb{Y}_{s,u}^{N-1} \otimes dX_u + \int_s^t \mathbb{Y}_{s,u}^{N-2} \otimes d\varphi_u\end{aligned}$$

because φ is antisymmetric. By induction hypothesis,

$$\begin{aligned}\mathbb{Y}^{N-1} &= \sum_k \sum_{\substack{\mathbf{I}=\mathbf{i_1}\ldots\mathbf{i_k}\in\{\mathbf{1},\mathbf{2}\}^k \\ i_1+\ldots+i_k=N-1}} a^{\mathbf{I}}, \\ \mathbb{Y}^{N-2} &= \sum_k \sum_{\substack{\mathbf{I}=\mathbf{i_1}\ldots\mathbf{i_k}\in\{\mathbf{1},\mathbf{2}\}^k \\ i_1+\ldots+i_k=N-2}} a^{\mathbf{I}}.\end{aligned}$$

Hence,

$$\mathbb{Y}^N = \sum_k \sum_{\substack{\mathbf{I}=\mathbf{i_1}\ldots\mathbf{i_k}\in\{\mathbf{1},\mathbf{2}\}^k \\ i_1+\ldots+i_k=N}} a^{\mathbf{I}},$$

as desired.

□

Lemma 3.2.11. Let X be a d-dimensional continuous semimartingale. Let $\ell \in T((\mathbb{R}^2)^*)$. Then, we have:

$$\int_0^T \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle dX_t = \langle \ell \mathbf{4}, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty} \rangle,$$

where the integral is in the sense of Itô, the notation $\ell\mathbf{4} \in T((\mathbb{R}^4)^*)$ means the concatenation of the word associated to ℓ with the letter $\mathbf{4}$ (introduced in Section

1.3) and $\widehat{\mathbb{X}}_{0,T}^{LL,<\infty}$ is the signature of the (4-dimensional) lead-lag process, as defined in Definition 2.3.1.

Proof. For a d-dimensional path Z, define the $2d$-dimensional path $\overline{Z} := (Z, Z)$, with the corresponding signature $\overline{\mathbb{Z}}^{<\infty}$. It was shown in [52] that $\widehat{\mathbb{X}}^{LL,<\infty}$ is the signature of the perturbed rough path $\widehat{\mathbb{X}}^{LL,\leq 2} := \overline{\widehat{\mathbb{X}}}^{\leq 2} + \psi$, with

$$\psi_{s,t} := \begin{pmatrix} 0 & -\frac{1}{2}[X]_{s,t} \\ \frac{1}{2}[X]_{s,t} & 0 \end{pmatrix}, \quad 0 \leq s \leq t \leq T.$$

Let $N \geq 1$. We have

$$\begin{aligned} \int_0^T \left\langle \ell, \widehat{\mathbb{X}}_{0,T}^{\leq N} \right\rangle dX_t &= \int_0^T \left\langle \ell, \widehat{\mathbb{X}}_{0,t}^{\leq N} \right\rangle \circ dX_t - \frac{1}{2} \left[\left\langle \ell, \int_0^{\cdot} \widehat{\mathbb{X}}_{0,u}^{\leq N-1} \otimes d\widehat{X}_u \right\rangle, X_{\cdot} \right]_T \\ &= \left\langle \ell \mathbf{4}, \overline{\widehat{\mathbb{X}}}_{0,T}^{\leq N+1} \right\rangle - \frac{1}{2} \left\langle \ell \mathbf{4}, \int_0^T \overline{\widehat{\mathbb{X}}}_{0,u}^{\leq N-1} \otimes d \left[\overline{\widehat{X}} \right]_u \right\rangle. \end{aligned}$$

Therefore, we have to show that

$$\left\langle \ell \mathbf{4}, \widehat{\mathbb{X}}_{0,T}^{\mathrm{LL},\leq N+1} \right\rangle = \left\langle \ell \mathbf{4}, \overline{\widehat{\mathbb{X}}}_{0,T}^{\leq N+1} - \frac{1}{2} \int_0^T \overline{\widehat{\mathbb{X}}}_{0,u}^{\leq N-1} \otimes d \left[\overline{\widehat{X}} \right]_u \right\rangle. \tag{A.1}$$

We proceed by induction. If $N = 1$, we have

$$\widehat{\mathbb{X}}_{0,T}^{\mathrm{LL},\leq 2} = \overline{\widehat{\mathbb{X}}}_{0,T}^{\leq 2} + \psi_{0,T}.$$

Since $\langle \ell \mathbf{4}, \psi \rangle = -\frac{1}{2} \left\langle \ell \mathbf{4}, \langle \overline{\widehat{X}} \rangle \right\rangle$, it follows that (A.1) holds for $N = 1$.

Assume that (A.1) holds for N, we will show that it also holds true for $N + 1$. By induction hypothesis, (A.1) is reduced to:

$$\langle \ell \mathbf{4}, \widehat{\mathbb{X}}_{0,T}^{\mathrm{LL},N+2} \rangle = \left\langle \ell \mathbf{4}, \overline{\widehat{\mathbb{X}}}_{0,T}^{N+2} - \frac{1}{2} \int_0^T \overline{\widehat{\mathbb{X}}}_{0,u}^{N} \otimes d \left[\overline{\widehat{X}} \right]_u \right\rangle.$$

By Lemma A.0.1,

$$\widehat{\mathbb{X}}_{0,T}^{\mathrm{LL},N+2} = \sum_{k} \sum_{\substack{\mathbf{I}=\mathbf{i_1}\dots\mathbf{i_k}\in\{\mathbf{1},\mathbf{2}\}^k \\ i_1+\dots+i_k=N+2}} a_{0,T}^{\mathbf{I}}.$$

Notice that $\langle \ell\mathbf{4}, a^{\mathbf{I}} \rangle$ is nonzero only for $\mathbf{I_1} = \underbrace{\mathbf{11\dots1}}_{N+2}$ and $\mathbf{I_2} = \underbrace{\mathbf{11\dots1}}_{N}\mathbf{2}$. Hence,

$$\begin{aligned}\langle \ell\mathbf{4}, \widehat{\mathbb{X}}_{0,T}^{\mathrm{LL},N+2} \rangle &= \langle \ell\mathbf{4}, a_{0,T}^{\mathbf{I_1}} \rangle + \langle \ell\mathbf{4}, a_{0,T}^{\mathbf{I_2}} \rangle = \left\langle \ell\mathbf{4}, \overline{\widehat{\mathbb{X}}}_{0,T}^{N+2} \right\rangle + \left\langle \ell\mathbf{4}, \int_0^T \overline{\widehat{\mathbb{X}}}_{0,u}^{N} \otimes d\psi_{0,u} \right\rangle \\ &= \left\langle \overline{\widehat{\mathbb{X}}}_{0,T}^{N+2} - \frac{1}{2}\int_0^T \overline{\widehat{\mathbb{X}}}_{0,u}^{N} \otimes d\left[\widehat{X}\right]_u \right\rangle\end{aligned}$$

as desired.

□

Theorem 3.3.3. Let $f \in T((\mathbb{R}^4)^*)$ and $p_0 \in \mathbb{R}$. Let $P \in \mathbb{R}[x]$ be a polynomial of one variable. Then, the solution of the optimal linear signature hedging problem (LSHP) is given by the solution of the following polynomial optimisation problem:

$$\inf_{\ell \in T((\mathbb{R}^2)^*)} \left\langle P^{\sqcup\!\sqcup}(f - p_0\varnothing - \ell\mathbf{4}), \mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\right] \right\rangle. \tag{A.2}$$

Proof. By Lemma 3.2.11, (LSHP) will be given by:

$$\begin{aligned}&\inf_{\ell \in T((\mathbb{R}^2)^*)} \mathbb{E}^{\mathbb{P}}\left[P\left(\langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle - p_0 - \int_0^T \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle dX_t\right)\right] \\ &= \inf_{\ell \in T((\mathbb{R}^2)^*)} \mathbb{E}^{\mathbb{P}}\left[P\left(\langle f, \widehat{\mathbb{X}}_{0,T}^{\mathrm{LL},<\infty}\rangle - p_0 - \langle \ell\mathbf{4}, \widehat{\mathbb{X}}_{0,T}^{\mathrm{LL},<\infty}\rangle\right)\right] \\ &= \inf_{\ell \in T((\mathbb{R}^2)^*)} \mathbb{E}^{\mathbb{P}}\left[P\left(\langle f - p_0\varnothing - \ell\mathbf{4}, \widehat{\mathbb{X}}_{0,T}^{\mathrm{LL},<\infty}\rangle\right)\right] \\ &\overset{(\star)}{=} \inf_{\ell \in T((\mathbb{R}^2)^*)} \mathbb{E}^{\mathbb{P}}\left[\langle P^{\sqcup\!\sqcup}\left(f - p_0\varnothing - \ell\mathbf{4}\right), \widehat{\mathbb{X}}_{0,T}^{\mathrm{LL},<\infty}\rangle\right] \\ &= \inf_{\ell \in T((\mathbb{R}^2)^*)} \left\langle P^{\sqcup\!\sqcup}\left(f - p_0\varnothing - \ell\mathbf{4}\right), \mathbb{E}^{\mathbb{P}}\left[\widehat{\mathbb{X}}_{0,T}^{\mathrm{LL},<\infty}\right] \right\rangle,\end{aligned}$$

where $(\star)$ follows by the shuffle product property (Definition 1.3.3).

□

Proposition 3.4.3. Let

$$a := \inf_{\theta \in \mathcal{T}^q(\Lambda_T)} \mathbb{E}\left[\exp\left(-\lambda\left(p_0 + \int_0^T \theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,t]})dX_t - F(\widehat{\mathbb{X}}^{LL,<\infty})\right)\right)\right]$$

be the infimum of the optimal exponential hedging problem. Given any $\varepsilon > 0$, there exists a polynomial $P_\varepsilon \in \mathbb{R}[x]$, a compact set $\mathcal{K}_\varepsilon \subset \widehat{\Omega}_T$, a linear signature payoff given by $f \in T((\mathbb{R}^4)^*)$ and a linear signature trading strategy given by $\ell \in T((\mathbb{R}^2)^*)$ such that:

1. $P_\varepsilon \xrightarrow{\varepsilon \to 0} \exp(-\lambda \cdot)$ uniformly on compacts,
2. $\mathbb{P}[\mathcal{K}_\varepsilon] > 1 - \varepsilon$,
3. $|F(\widehat{\mathbb{X}}^{LL,<\infty}) - \langle f, \widehat{\mathbb{X}}^{LL,<\infty}_{0,T}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}} \in \mathcal{K}_\varepsilon$,
4. $|\theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,t]}) - \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}^{<\infty} \in \mathcal{K}_\varepsilon$ and $t \in [0,T]$,
5. $|a_\varepsilon - a| \leq \varepsilon$, where

$$a_\varepsilon := \mathbb{E}\left[P_\varepsilon\left(p_0 + \int_0^T \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle dX_t - \langle f, \widehat{\mathbb{X}}^{LL,<\infty}_{0,T}\rangle\right) \;;\; \mathcal{K}_\varepsilon\right].$$

Proof. Let $I \subset \mathbb{R}$ be a compact interval such that $p_0 + \int_0^T \theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,t]})dX_t - F(\widehat{\mathbb{X}}^{LL,<\infty}) \in I$ a.s. Let P_ε be the Taylor expansion of $\exp(-\lambda\,\cdot)$ around the origin of degree large enough so that $\|P_\varepsilon - \exp(-\lambda\,\cdot)\|_{L^\infty(I)} < \varepsilon$.

Let $\mathcal{K}_\varepsilon \subset \widehat{\Omega}_p$ compact be such that $\mathbb{P}[\mathcal{K}_\varepsilon] > 1 - \varepsilon$ and

$$\left|\mathbb{E}^{\mathbb{P}}\left[P_\varepsilon\left(p_0 + \int_0^T \theta(\widehat{\mathbb{X}}|^{<\infty}_{[0,t]})dX_t - F(\widehat{\mathbb{X}}^{LL,<\infty})\right) \;;\; \mathcal{K}_\varepsilon^c\right]\right| < \varepsilon.$$

Take $\ell \in T((\mathbb{R}^2)^*)$, $f \in T((\mathbb{R}^4)^*)$ such that 3. and 4. hold, which we can do due to Proposition 3.3.7 and Proposition 3.3.8. Then, we have:

$$
\begin{aligned}
&\left|\mathbb{E}\left[P_\varepsilon\left(p_0+\int_0^T\langle\ell,\widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle dX_t-\langle f,\widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle\right)-\exp\left(-\lambda\left(p_0+\int_0^T\theta_t dX_t-F(\widehat{\mathbb{X}}^{LL,<\infty})\right)\right)\right]\right|\\
&\leq\left|\mathbb{E}\left[P_\varepsilon\left(p_0+\int_0^T\langle\ell,\widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle dX_t-\langle f,\widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle\right)-P_\varepsilon\left(p_0+\int_0^T\theta_t dX_t-F(\widehat{\mathbb{X}}^{LL,<\infty})\right)\right]\right|\\
&+\left|\mathbb{E}\left[P_\varepsilon\left(p_0+\int_0^T\theta_t dX_t-F(\widehat{\mathbb{X}}^{LL,<\infty})\right)-\exp\left(-\lambda\left(p_0+\int_0^T\theta_t dX_t-F(\widehat{\mathbb{X}}^{LL,<\infty})\right)\right)\right]\right|\\
&=:(\star)+(\star\star)
\end{aligned}
$$

By Propositions 3.3.7 and 3.3.8, $(\star) < \varepsilon$. Moreover, because $\|P_\varepsilon - \exp(-\lambda\,\cdot)\|_{L^\infty(I)} < \varepsilon$, we have $(\star\star) < \varepsilon$, and the proof follows. □

Corollary 3.4.7. The solution of the optimal hedging problem under fixed quadratic trading costs (3.3) is given by the solution of the following optimisation problem:

$$
\inf_{v\in T((\mathbb{R}^2)^*)}\left\langle P^{\text{ш}}(f-p_0\varnothing-v\mathbf{14}+\alpha v^{\text{ш}2}\mathbf{1}),\mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\right]\right\rangle. \tag{A.3}
$$

Similarly, the solution of the optimal hedging under proportional transaction costs (3.4) is given by

$$
\inf_{v\in T((\mathbb{R}^2)^*)}\left\langle P^{\text{ш}}(f-p_0\varnothing-v\mathbf{14}+\alpha(v\text{ш}(\mathbf{2}+\varnothing))^{\text{ш}2}\mathbf{1}),\mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\right]\right\rangle. \tag{A.4}
$$

Proof. (A.3) follows from Theorem 3.3.3 and from the fact that

$$
\int_0^T|\langle v,\widehat{\mathbb{X}}_{0,u}^{<\infty}\rangle|^2du=\langle v^{\text{ш}2}\mathbf{1}\widehat{\mathbb{X}}_{0,u}^{<\infty}\rangle.
$$

Similarly, (A.4) follows from Theorem 3.3.3 and the fact that

$$
\int_0^T|\langle v,\widehat{\mathbb{X}}_{0,u}^{<\infty}\rangle X_u|^2du=\langle(v\text{ш}(\mathbf{2}+\varnothing))^{\text{ш}2}\mathbf{1}\widehat{\mathbb{X}}_{0,u}^{<\infty}\rangle.
$$

□

Appendix B

From the linear signature hedging problem to the polynomial hedging problem

In Section 3.3.2, we justified why solving the linear signature hedging problem (LSHP) allows us to numerically estimate the solution of the polynomial hedging problem (PHP). The objective of this section is to give the technical details of why such an approximation is justified.

Proposition 3.3.6. Let $F : \widehat{\Omega}_T^{LL} \to \mathbb{R}$ be a continuous payoff and let $\mathcal{K} \subset \widehat{\Omega}_T^{LL}$ be a compact set. Given any $\varepsilon > 0$, there exists $f \in T((\mathbb{R}^4)^*)$ such that

$$|F(\widehat{\mathbb{X}}^{LL,<\infty}) - \langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}^{LL,<\infty} \in \mathcal{K}.$$

Proof. Let $f, g \in T((\mathbb{R}^4)^*)$. Then, by the shuffle product, we have:

$$\langle f, \widehat{\mathbb{X}}_{0,T}^{<\infty}\rangle\langle g, \widehat{\mathbb{X}}_{0,T}^{<\infty}\rangle = \langle f \sqcup\!\sqcup g, \widehat{\mathbb{X}}_{0,T}^{<\infty}\rangle \quad \forall\, \mathbb{X}^{LL,<\infty} \in \widehat{\Omega}_T.$$

Therefore, linear signature payoffs form an algebra. Moreover, the uniqueness of the signature implies that the family of linear signature payoffs separate points. Also, they trivially contain constants.

Let $\varepsilon > 0$ and let $\mathcal{K} \subset \widehat{\Omega}_T^{LL}$ be a compact set. Applying the Stone–Weierstrass theorem, we deduce there exists an f-linear signature payoff with $f \in T((\mathbb{R}^4)^*)$ such

that

$$|F(\widehat{\mathbb{X}}^{LL,<\infty}) - \langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}^{LL,<\infty} \in \mathcal{K}.$$

□

Corollary 3.3.7. Let $F : \widehat{\Omega}_T^{LL} \to \mathbb{R}$ be a continuous payoff. Given any $\varepsilon > 0$, there exists a compact set $\mathcal{K}_\varepsilon \subset \widehat{\Omega}_T$ (which does not depend on F) and $f \in T((\mathbb{R}^4)^*)$ such that:

1. $\mathbb{P}[\mathcal{K}_\varepsilon] > 1 - \varepsilon$,
2. $|F(\widehat{\mathbb{X}}^{LL,<\infty}) - \langle f, \widehat{\mathbb{X}}_{0,T}^{<\infty}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}^{LL,<\infty} \in \mathcal{K}_\varepsilon$.

Proof. Let $\varepsilon > 0$. Given that $\widehat{\Omega}_T$ is separable, there exists $\mathcal{K}_\varepsilon \subset \widehat{\Omega}_T$ compact such that $\mathbb{P}[\mathcal{K}_\varepsilon] > 1 - \varepsilon$. By the previous proposition, there exists an f-linear signature payoff, with $f \in T((\mathbb{R}^4)^*)$, such that $|F(\widehat{\mathbb{X}}^{LL,<\infty}) - \langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}^{LL,<\infty} \in \mathcal{K}_\varepsilon$. □

Proposition B.0.1. *Set $1 \leq q < \infty$, and let $F : \widehat{\Omega}_T^{LL} \to \mathbb{R}$ be an L^q-payoff. Given any $\varepsilon > 0$, there exists a compact set $\mathcal{K}_\varepsilon \subset \widehat{\Omega}_p$ and a f-linear signature payoff with $f \in T((\mathbb{R}^4)^*)$ such that:*

1. $\mathbb{P}[\mathcal{K}_\varepsilon] > 1 - \varepsilon$,
2. $\mathbb{E}[|F(\widehat{\mathbb{X}}^{LL,<\infty}) - \langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle|^q \,;\, \mathcal{K}_\varepsilon] < \varepsilon$.

Proof. Let $\varepsilon > 0$. Since continuous functions are dense in L^q, there exists a continuous payoff $G : \widehat{\Omega}_T \to \mathbb{R}$ such that $\|F - G\|_{L^q}^q < \varepsilon/2$. By Proposition B.0.1 we may pick $\mathcal{K}_\varepsilon \subset \widehat{\Omega}$ and $f \in T((\mathbb{R}^4)^*)$ such that $\mathbb{P}[\mathcal{K}_\varepsilon] > 1 - \varepsilon$, $\mathbb{E}[|G(\widehat{\mathbb{X}}^{LL,<\infty})|^q; \mathcal{K}_\varepsilon^c] < \varepsilon$ and

$$|G(\widehat{\mathbb{X}}^{<\infty}) - \langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle|^q < \varepsilon/2 \quad \forall\, \widehat{\mathbb{X}}^{LL,<\infty} \in \mathcal{K}_\varepsilon.$$

Then,

$$\begin{aligned}\mathbb{E}[|F(\widehat{\mathbb{X}}^{LL,<\infty}) - \langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle|^q \,;\, \mathcal{K}_\varepsilon] &\leq \mathbb{E}[|F(\widehat{\mathbb{X}}^{LL,<\infty}) - G(\widehat{\mathbb{X}}^{LL,<\infty})|^q \,;\, \mathcal{K}_\varepsilon] \\ &\quad + \mathbb{E}[|G(\widehat{\mathbb{X}}^{LL,<\infty}) - \langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle|^q \,;\, \mathcal{K}_\varepsilon] \\ &< \varepsilon/2 + \varepsilon/2 = \varepsilon.\end{aligned}$$

□

Proposition 3.3.8. Let $\mathcal{K} \subset \Lambda_T$ be a compact set. Then, given any trading strategy $\theta \in \mathcal{T}$, there exists $\ell \in T((\mathbb{R}^2)^*)$ such that

$$|\theta(\widehat{\mathbb{X}}|_{[0,t]}^{<\infty}) - \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}|_{[0,t]}^{<\infty} \in \mathcal{K}.$$

Proof. Given $\theta_1, \theta_2 \in \mathcal{T}_{\text{sig}}(\mathcal{K})$, there exist $\ell_1, \ell_2 \in T((\mathbb{R}^2)^*)$ such that $\theta_i(\widehat{\mathbb{X}}|_{[0,t]}) = \langle \ell_i, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle$ for each $\widehat{\mathbb{X}}|_{[0,t]} \in \mathcal{K}$ and $i = 1, 2$. Define $\theta_{1,2}(\mathbb{X}|_{[0,t]}) := \langle \ell_1 \text{ш} \ell_2, \mathbb{X}_{0,t}^{<\infty}\rangle$. Then,

$$\begin{aligned}\theta_1(\widehat{\mathbb{X}}|_{[0,t]})\theta_2(\widehat{\mathbb{X}}|_{[0,t]}) &= \langle \ell_1, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle\langle \ell_2, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle \\ &= \langle \ell_1 \text{ш} \ell_2, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle \\ &= \theta_{1,2}(\widehat{\mathbb{X}}|_{[0,t]}).\end{aligned}$$

Hence, $\mathcal{T}_{\text{sig}}(\mathcal{K})$ is an algebra. Given that it also separates points ([15]) and contains constants, it follows by Stone–Weierstrass theorem that given any trading strategy $\theta \in \mathcal{T}$, there exists $\ell \in T((\mathbb{R}^2)^*)$ such that

$$|\theta(\widehat{\mathbb{X}}|_{[0,t]}^{<\infty}) - \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}|_{[0,t]}^{<\infty} \in \mathcal{K}.$$

□

Theorem B.0.2. *Let*

$$a := \inf_{\theta \in \mathcal{T}^q(\Lambda_T)} \mathbb{E}\left[P\left(F(\widehat{\mathbb{X}}^{LL,<\infty}) - p_0 - \int_0^T \theta(\widehat{\mathbb{X}}|_{[0,t]}^{<\infty}) dX_t\right)\right]$$

be the infimum of the optimal polynomial hedging problem (PHP)*. Given any $\varepsilon > 0$, there exists a compact set $\mathcal{K}_\varepsilon \subset \widehat{\Omega}_T$, a linear signature payoff given by $f \in T((\mathbb{R}^4)^*)$ and a linear signature trading strategy given by $\ell \in T((\mathbb{R}^2)^*)$ such that:*

1. $\mathbb{P}[\mathcal{K}_\varepsilon] > 1 - \varepsilon$,

2. $|F(\widehat{\mathbb{X}}^{LL,<\infty}) - \langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}^{LL,<\infty} \in \mathcal{K}_\varepsilon$,

3. $|\theta(\widehat{\mathbb{X}}|_{[0,t]}^{<\infty}) - \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle| < \varepsilon \quad \forall\, \widehat{\mathbb{X}}^{<\infty} \in \mathcal{K}_\varepsilon$ *and* $t \in [0,T]$,

4. $|a_\varepsilon - a| \leq \varepsilon$, *where*

$$a_\varepsilon := \mathbb{E}\left[P\left(\langle f, \widehat{\mathbb{X}}_{0,T}^{LL,<\infty}\rangle - p_0 - \int_0^T \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty}\rangle dX_t\right) \;;\; \mathcal{K}_\varepsilon\right].$$

Proof. Follows from Proposition B.0.1 and Proposition 3.3.8, and the triangle inequality. □

Appendix C

Proofs of Chapter 4

Proof of Proposition 4.3.2. First, we show that the class of all signature payoff forms an algebra. Let $\ell_1, \ell_2 \in T((\mathbb{R}^2 \oplus \mathbb{R})^*)$, and consider the corresponding signature payoffs $\mathcal{S}^{\ell_1}, \mathcal{S}^{\ell_2} : \widehat{\Omega} \to \mathbb{R}$. It is clear that $\mathcal{S}^{\ell_1} + \mathcal{S}^{\ell_2} = \mathcal{S}^{\ell_1+\ell_2}$ – in other words, the sum of two signature payoffs is a signature payoff. Moreover, by the shuffle product property, we have:

$$\mathcal{S}^{\ell_1} \cdot \mathcal{S}^{\ell_2} = \mathcal{S}^{\ell_1 \text{ш} \ell_2}.$$

That is, the product of two signature payoffs is another signature payoff, and the linear functional is given by the shuffle product of the linear functionals of the original signature payoffs. Therefore, the space of signature payoffs forms an algebra.

By Lemma 4.2.2, the signature map $\widehat{\Omega} \to T((\mathbb{R}^2 \oplus \mathbb{R}))$ is injective. Therefore, given $\widehat{X}_1, \widehat{X}_2 \in \widehat{\Omega}$ distinct, we have $S(\widehat{X}_1) \neq S(\widehat{X}_2)$. It follows immediately that there exists a signature payoff $\mathcal{S}^\ell$, with $\ell \in T((\mathbb{R}^2 \oplus \mathbb{R})^*)$, such that $\mathcal{S}^\ell(\widehat{X}_1) \neq \mathcal{S}^\ell(\widehat{X}_2)$. Therefore, signature payoffs separate points.

Given that the space of signature payoffs trivially contain constants, we conclude by the Stone–Weierstrass theorem that given any continuous payoff $G : \widehat{\Omega} \to \mathbb{R}$ and a compact set $\mathcal{K} \subset \widehat{\Omega}$, there exists a signature payoff $\mathcal{S}^\ell$ with $\ell \in T((\mathbb{R}^2 \oplus \mathbb{R})^*)$ such that $|G(\widehat{X}) - \mathcal{S}^\ell(\widehat{X})| < \varepsilon$ for all $\widehat{X} \in \mathcal{K}$. □

Proof of Corollary 4.3.3. By tightness of $\mathcal{M}$, we pick a compact set $\mathcal{K}_\varepsilon \subset \widehat{\Omega}$ such that

$\mathbb{Q}(\widehat{X} \in \mathcal{K}_\varepsilon) > 1 - \varepsilon$ for all risk-neutral measures $\mathbb{Q} \in \mathcal{M}$. Apply Proposition 4.3.2 with the payoff G and compact $\mathcal{K}_\varepsilon$ and conclude that $\|G - \mathcal{S}^\ell\|_{L^\infty(\mathcal{K}_\varepsilon)} < \varepsilon$. □

Proof of Proposition 4.3.4. Substitute G with a continuous payoff that is close to G in L^1 if necessary and assume, without loss of generality, that G is continuous.

By Corollary 4.3.3, choose a compact set $\mathcal{K}_1 \subset \widehat{\Omega}$ such that $\mathbb{Q}(\widehat{X} \in \mathcal{K}_1) > 1 - \varepsilon$ for all $\mathbb{Q} \in \mathcal{M}$ and $\|G - \mathcal{S}^\ell\|_{L^\infty(\mathcal{K}_1)} < \varepsilon/2$.

Moreover, pick $\mathcal{K}_2 \subset \widehat{\Omega}$ compact such that $\mathbb{E}^{\mathbb{Q}}[|G(\widehat{X})| \; ; \; \mathcal{K}_2^c|] < \varepsilon/2$. Indeed:

1. If G is bounded, then by tightness of $\mathcal{M}$ we may pick a compact set $\mathcal{K}_2 \subset \widehat{\Omega}$ such that $\mathbb{E}^{\mathbb{Q}}[|G(\widehat{X})| \; ; \; \mathcal{K}_2^c|] < \varepsilon/2$ for all $\mathbb{Q} \in \mathcal{M}$.

2. Otherwise, if $G \in \mathcal{A}_{\varepsilon/4}$, there exists by definition an attainable payoff $H \in \mathcal{A}$ such that $\|H - G\|_{L^\infty(\widehat{\Omega})} < \varepsilon/4$. Let $\mathcal{K}_2 \subset \widehat{\Omega}$ be a compact set such that $\mathbb{E}^{\mathbb{Q}}[|H(\widehat{X})| \; ; \; \mathcal{K}_2^c|] < \varepsilon/4$ for all $\mathbb{Q} \in \mathcal{M}$. Then, $\mathbb{E}^{\mathbb{Q}}[|G(\widehat{X})| \; ; \; \mathcal{K}_2^c|] < \varepsilon/2$ for all $\mathbb{Q} \in \mathcal{M}$.

Set $\mathcal{K}_\varepsilon := \mathcal{K}_\varepsilon$. Then, $\mathbb{Q}(\widehat{X} \in \mathcal{K}_\varepsilon) > 1 - \varepsilon$ and

$$\begin{aligned}\mathcal{P}(|G - \mathcal{L}|) &= \sup_{\mathbb{Q}\in\mathcal{M}} \mathbb{E}^{\mathbb{Q}}[|G(\widehat{X}) - \mathcal{L}(\widehat{X})] = \sup_{\mathbb{Q}\in\mathcal{M}} (\mathbb{E}^{\mathbb{Q}}[|G(\widehat{X})| \; ; \; \mathcal{K}_\varepsilon^c] + \mathbb{E}^{\mathbb{Q}}[|G(\widehat{X}) - \mathcal{S}^\ell(\widehat{X})| \; ; \; \mathcal{K}_\varepsilon]) \\ &< \varepsilon.\end{aligned}$$

□

Proof of 4.4.1. We will show that the space of signature trading strategies is a subalgebra of the space of trading strategies $\mathcal{T}$. Let $\ell_1, \ell_2 \in T((\mathbb{R}^2)^*)$. Then, by the shuffle product property, we have:

$$\langle \ell_1, \mathbb{X}_{0,t}^{b,<\infty} \rangle \langle \ell_1, \mathbb{X}_{0,t}^{b,<\infty} \rangle = \langle \ell_1 ш \ell_2, \mathbb{X}_{0,t}^{b,<\infty} \rangle \quad \forall \, \widehat{X}|_{[0,t]} \in \Lambda.$$

In other words, the product of two signature trading strategies is another signature strategy. On the other hand, it is trivial that the sum of two signature trading strategies is also a signature trading strategy. Therefore, the space of signature trading strategies is a subalgebra of $\mathcal{T}$.

Moreover, given $\widehat{X}|_{[0,t]}, \widehat{Y}|_{[0,s]} \in \Lambda$ with $\widehat{X}|_{[0,t]} \neq \widehat{Y}|_{[0,s]}$, we claim there exists a linear functional $\ell \in T((\mathbb{R}^2)^*)$ such that $\langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \neq \langle \ell, \widehat{\mathbb{Y}}_{0,s}^{<\infty} \rangle$ – in other words, we claim that the space of signature trading strategies separates points. If $t \neq s$, set $\ell := \mathbf{1} \in T((\mathbb{R}^2)^*)$. Then,

$$\langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle = t \neq s = \langle \ell, \widehat{\mathbb{Y}}_{0,s}^{<\infty} \rangle.$$

Assume, on the other hand, that $t = s$. By Lemma 4.2.2, we have $\widehat{\mathbb{X}}_{0,t}^{<\infty} \neq \widehat{\mathbb{Y}}_{0,s}^{<\infty}$. Therefore, there exists $\ell \in T((\mathbb{R}^2)^*)$ such that

$$\langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \neq \langle \ell, \widehat{\mathbb{Y}}_{0,s}^{<\infty} \rangle,$$

thus proving our claim.

Let $\varepsilon > 0$. Pick $\mathcal{K}_\varepsilon \subset \widehat{\Omega}$ such that $\mathbb{Q}(\widehat{X} \in \mathcal{K}_\varepsilon) > 1 - \varepsilon$ for all $\mathbb{Q} \in \mathcal{M}$. As the space of signature trading strategies forms an algebra, separates points and it trivially contains constants, we conclude by Stone–Weierstrass theorem that given $\theta \in \mathcal{T}$, there exists $\ell \in T((\mathbb{R}^2)^*)$ such that

$$|\theta(X^b|_{[0,t]}) - \langle \ell, \mathbb{X}_{0,t}^{b,<\infty} \rangle| < \varepsilon \quad \forall\, \widehat{X} \in \mathcal{K}_\varepsilon,\, t \in [0,1].$$

□

Appendix D

Proofs of Chapter 6

We assume the price impacts are given, for a signature trading speed $\boldsymbol{\theta} = (\theta^1, \theta^2)$, by

$$P_t^{\theta^i} = P_t^i - \langle g^{\theta^i}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \,, \tag{D.1}$$

$$\widetilde{P}_t^{\theta^i} = P_t^{\theta^i} - \langle h^{\theta^i}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \,, \tag{D.2}$$

for $i = 1, 2$, where $g^{\theta^i} \in T((\mathbb{R}^3)^*)$ is the temporary market impact and $h^{\theta^i} \in T((\mathbb{R}^3)^*)$ is the permanent market impact. We note that the framework of Chapter 6 is a particular case of this setting.

Lemma D.0.1. *Let $P_0^1 = s^1$ and $P_0^2 = s^2$. Let $\boldsymbol{\theta} = (\theta^1, \theta^2)$ be signature trading speeds in the stock and the currency pair, given by $\theta_t^1 = \langle \beta, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$ and $\theta_t^2 = \langle \gamma, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$, with $\beta, \gamma \in T((\mathbb{R}^3)^*)$. Then,*

1. $P_t^\beta = \langle p^\beta, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \,,$
2. $P_t^\gamma = \langle p^\gamma, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \,,$
3. $\widetilde{P}_t^\beta = \langle \tilde{p}^\beta, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \,,$
4. $\widetilde{P}_t^\gamma = \langle \tilde{p}^\gamma, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \,,$
5. $W_t^\gamma = \left\langle (\tilde{p}^\gamma \sqcup\!\sqcup \gamma)\, \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle = \langle w^\gamma, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \,,.$
6. $Q_t^\beta = \langle q_0 \varnothing - \beta \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle = \langle q^\beta, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \,,$

7. $Q_t^\gamma = \langle q^\gamma + w^\beta, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$,

8. $\int_0^t \left(Q_s^\beta \, P_s^\beta \, P_s^\gamma\right)^2 \mathrm{d}s = \left\langle \left(q^\beta \sqcup\!\sqcup p^\beta \sqcup\!\sqcup p^\gamma\right)^{\sqcup\!\sqcup 2} \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle$,

9. $\int_0^t \left(Q_s^\gamma \, P_s^\gamma\right)^2 \mathrm{d}s = \left\langle \left(\left(q^\gamma + w^\beta\right) \sqcup\!\sqcup p^\gamma\right)^{\sqcup\!\sqcup 2} \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle$,

10. $Q_t^\beta \, P_t^\beta \, \left(P_t^\gamma - \alpha_1 Q_t^\beta \, P_t^\beta\right) = \left\langle m^\beta \sqcup\!\sqcup \left(p^\gamma - \alpha_1 \, m^\beta\right), \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle$,

11. $Q_t^\gamma \left(P_t^\gamma - \alpha_2 Q_t^\gamma\right) = \left\langle \left(q^\gamma + w^\beta\right) \sqcup\!\sqcup \left(p^\gamma - \alpha_2 \, q^\gamma - \alpha_2 \, w^\beta\right), \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle$,

with

$$
\begin{array}{ll}
p^\beta = \mathbf{2} + s^1 \varnothing - g^\beta \,, & p^\gamma = \mathbf{3} + s^2 \varnothing - g^\gamma \,, \\
\tilde{p}^\beta = \mathbf{2} + s^1 \varnothing - \left(g^\beta + h^\beta\right) \,, & \tilde{p}^\gamma = \mathbf{3} + s^2 \varnothing - \left(g^\gamma + h^\gamma\right) \,, \\
w^\beta = \left[\tilde{p}^\beta \sqcup\!\sqcup \beta\right] \mathbf{1} \,, & w^\gamma = \left[\tilde{p}^\gamma \sqcup\!\sqcup \gamma\right] \mathbf{1} \,, \\
q^\beta = q_0 \varnothing - \beta \mathbf{1} \,, & q^\gamma = -\gamma \mathbf{1} \,, \\
m^\beta = q^\beta \sqcup\!\sqcup p^\beta \,. &
\end{array}
$$

Proof. Let $t \in [0, T]$.

1. Notice that because $P_0^1 = s^1$, we have $P_u^1 = \langle \mathbf{2} + s^1 \varnothing, \widehat{\mathbb{X}}_{0,u}^{<\infty} \rangle$ for each $u \in [0, t]$. Then $P_u^\beta = \langle \mathbf{2} + s^1 \varnothing - g^\beta, \widehat{\mathbb{X}}_{0,u}^{<\infty} \rangle = \langle p^\beta, \widehat{\mathbb{X}}_{0,u}^{<\infty} \rangle$.

2. Similar to previous proof.

3. Similar to previous proof.

4. Similar to previous proof.

5. Then, by the shuffle product property,

$$
\begin{aligned}
W_t^\gamma &= \int_0^t \widetilde{P}_s^\gamma \, \langle \gamma, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle \, \mathrm{d}s = \int_0^t \langle \tilde{p}^\beta, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle \, \langle \gamma, \widehat{\mathbb{X}}_{0,s}^{<\infty} \rangle \, \mathrm{d}s \\
&= \int_0^t \left\langle \tilde{p}^\gamma \sqcup\!\sqcup \gamma, \widehat{\mathbb{X}}_{0,s}^{<\infty} \right\rangle \mathrm{d}s = \left\langle \left(\tilde{p}^\gamma \sqcup\!\sqcup \gamma\right) \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle = \langle w^\gamma, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle \,.
\end{aligned}
$$

6. From the definition of Q^β we have

$$
\begin{aligned}
Q_t^\beta &= q_0 - \int_0^t \left\langle \beta, \widehat{\mathbb{X}}_{0,s}^{<\infty} \right\rangle \mathrm{d}s \\
&= \left\langle q_0 \varnothing, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle - \left\langle \beta \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle = \left\langle q^\beta, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle \,.
\end{aligned}
$$

7. From the definition of Q^{θ^2} we have

$$\begin{aligned} Q_t^{\theta^2} &= -\int_0^t \left\langle \gamma, \widehat{\mathbb{X}}_{0,s}^{<\infty} \right\rangle \mathrm{d}s + \int_0^t \left\langle \beta, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle \left\langle \tilde{p}^\beta, \widehat{\mathbb{X}}_{0,s}^{<\infty} \right\rangle \mathrm{d}s \\ &= \left\langle -\gamma \mathbf{1} + w^\beta, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle = \left\langle q^\gamma + w^\beta, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle . \end{aligned}$$

8. By the shuffle product property we have

$$\begin{aligned} \int_0^t \left(Q_s^\beta \, P_s \, P_s^\gamma\right)^2 \mathrm{d}s &= \int_0^t \left\langle q^\beta \sqcup\!\sqcup p^\beta \sqcup\!\sqcup p^\gamma, \widehat{\mathbb{X}}_{0,s}^{<\infty} \right\rangle \left\langle q^\beta \sqcup\!\sqcup p^\beta \sqcup\!\sqcup p^\gamma, \widehat{\mathbb{X}}_{0,s}^{<\infty} \right\rangle \mathrm{d}s \\ &= \left\langle \left(q^\beta \sqcup\!\sqcup p^\beta \sqcup\!\sqcup p^\gamma\right)^{\sqcup\!\sqcup 2} \mathbf{1}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle . \end{aligned}$$

9. Similar to previous proof.

10. Follows from the shuffle product property.

11. Follows from the shuffle product property.

□

Therefore, the double-execution problem (5.3) is written as:

Proposition D.0.2. *Let $P_0^1 = s^1$ and $P_0^2 = s^2$. Let $\boldsymbol{\theta} = (\theta^1, \theta^2)$ be a signature trading speed, given by $\theta_t^1 = \langle \beta, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$ and $\theta_t^2 = \langle \gamma, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$, with $\beta, \gamma \in T((\mathbb{R}^3)^*)$. Then, we have*

$$\begin{aligned} V^{\boldsymbol{\theta}} = \Big\langle & w^\gamma + m^\beta \sqcup\!\sqcup \left(p^\gamma - \alpha_1 \, m^\beta\right) + \left(q^\gamma + w^\beta\right) \sqcup\!\sqcup \left(p^\gamma - \alpha_2 \, q^\gamma - \alpha_2 \, w^\beta\right) && \text{(D.3)} \\ & - \phi_1 \, \left(q^\beta \sqcup\!\sqcup p^\beta \sqcup\!\sqcup p^\gamma\right)^{\sqcup\!\sqcup 2} \mathbf{1} - \phi_2 \, \left(\left(q^\gamma + w^\beta\right) \sqcup\!\sqcup p^\gamma\right)^{\sqcup\!\sqcup 2} \mathbf{1}, \widehat{\mathbb{X}}_{0,T}^{<\infty} \Big\rangle . && \text{(D.4)} \end{aligned}$$

Therefore, the double-execution problem (5.3) *reduces to*

$$\begin{aligned} \sup_{\beta,\gamma \in T((\mathbb{R}^3)^*)} \Big\langle & w^\gamma + m^\beta \sqcup\!\sqcup \left(p^\gamma - \alpha_1 \, m^\beta\right) + \left(q^\gamma + w^\beta\right) \sqcup\!\sqcup \left(p^\gamma - \alpha_2 \, q^\gamma - \alpha_2 \, w^\beta\right) && \text{(D.5)} \\ & - \phi_1 \, \left(q^\beta \sqcup\!\sqcup p^\beta \sqcup\!\sqcup p^\gamma\right)^{\sqcup\!\sqcup 2} \mathbf{1} - \phi_2 \, \left(\left(q^\gamma + w^\beta\right) \sqcup\!\sqcup p^\gamma\right)^{\sqcup\!\sqcup 2} \mathbf{1}, \mathbb{E}\left[\widehat{\mathbb{X}}_{0,T}^{<\infty}\right] \Big\rangle . \end{aligned}$$

Proof of Proposition D.0.2. Follows from Lemma D.0.1. □

Appendix E

Deep Signature Transforms: Implementation Details

All experimental models were trained using the Adam [89] optimiser as implemented by PyTorch [115], which was the framework used to implement the models. Signature calculations were performed with the aid of the `iisignature` package [119]. All activation functions were taken to be the ReLU. Computations were performed on two computers. One was equipped with two Tesla K40m GPUs. The second was equipped with two GeForce RTX 2080 Ti GPUs and two Quadro GP100 GPUs.

In each of the following subsections, the notation is the same as the notation used in the corresponding section of the main document.

E.1 A generative model for a stochastic process

The training dataset was given by 1024 realisations of an Ornstein–Uhlenbeck process, and the test set was of the same size, each sampled at 100 points of $[0, 1]$. No minibatching was used. The model was trained for 500 epochs.

The layer Φ^{θ_1} operated pointwise on the stream of time-augmented Brownian motion $\mathbf{B} = (t_i, B_{t_i})_i \in \mathcal{S}(\mathbb{R}^2)$, and was taken to be a neural network with 2 output neurons and 2 hidden layers of 8 neurons. Furthermore it kept the original stream;

thus

$$\Phi^{\theta_1}(\mathbf{B}) = (t_i, B_{t_i}, \phi_1^{\theta_1}(t_i, B_{t_i}), \phi_2^{\theta_1}(t_i, B_{t_i})) \in \mathcal{S}(\mathbb{R}^4)$$

for some learned $\phi_1^{\theta_1}, \phi_2^{\theta_1}$. The lift was the 'expanding window' described in equation (7.5). The signature in the generator was truncated at $N = 3$ (giving 84 scalar nonconstant terms) The layer f^{θ_2} operated pointwise on the stream of signatures, and was a simple linear map down to a scalar value (the value of the generated process at that time step). The signature in the discriminator was truncated at $M = 4$.

Some hyperparameter searching was necessary to obtain good results. The search was not done according to any formal scheme. It seemed that if Φ^{θ_1} was sufficiently simple and did not keep the original stream then the training would easily get trapped in a bad local minima, and the generated process would be visually distinct from the Ornstein–Uhlenbeck process.

E.2 Supervised learning with fractional Brownian motion

The training set featured 600 samples whilst the test set featured 100 samples, each of an instance of fractional Brownian motion sampled at 300 time steps of $[0, 1]$, with Hurst parameters in the range $[0.2, 0.8]$. These were split up into batches of 128 samples, so the last batch is slightly smaller than the others, and every model trained for 100 epochs. The loss function was taken to be mean squared error (MSE).

There was no hyperparameter searching except to require that all models should have approximately the same number of parameters; in all cases the results represent a model whose hyperparameters have not been fine-tuned to the task at hand.

All models used a sigmoid as a final nonlinearity, so as to map in to $(0, 1)$.

The differing sizes of layers between models (whilst keeping roughly the same overall parameter count) is usually because of the varying size of the input to the

model. Some models take all of the raw data, some models use signatures, and some models take expanding or sliding windows of the data in a manner akin to equations (7.5) and (7.7).

The Feedforward model was a simple neural network with 3 hidden layers of 16 neurons each.

The Neural-Sig model – which is essentially the same model as the Feedforward model, except that the data has the signature applied as feature transformation first – featured hidden layers of sizes 64, 64, 32, 32, 16, 16 respectively.

The RNN model is two recurrent neural networks, the first comprised of dense layers of sizes 64, 64, 32 and output size 6, and the second comprised of dense layers of size 32, 32, 32, and output size 5. The first network sweeps across the input data relatively slowly, with a stride of 2, whilst the second network sweeps across the result of the first network more quickly, with a stride of 4. In this way it may capture information from the input data at multiple timescales; part of the challenge of fractional Brownian motion is the existence of long-range dependencies [51].

The LSTM and GRU models both featured two recurrent layers each of size 32, and swept across the raw data with a stride of 1.

DeepSigNet featured a single Neural-Lift-Signature block, where the neural component was given by a single convolutional layer with 3 channels and kernel size 3, the lift was the trivial lift of equation (7.8), and the signature was truncated as $N = 3$. The neural component also preserved the original time-augmented stream of data, so that in some sense the neural component has 3 extra channels corresponding to time and value. On top of this a feedforward neural network with 5 hidden layers of size 32 was placed. Thus this model is very similar to the Neural-Sig model, except that a small learnable transformation was allowed before the signature. The difference in their performance highlights the value of learning a transformation before using the

signature. (Without which the Neural-Sig model is merely outperformed by some non-signature based models.)

DeeperSigNet featured three Neural-Lift-Signature blocks. The neural component of the first block was a small feedforward network with 2 hidden layers of size 16 and an output layer of size 3, swept across the length of the stream; its kernel size (how many time-value pairs of the stream it saw at once) was 4. The original time-augmented stream of data was also preserved by the neural component. The neural components of the other two blocks were recurrent neural networks, featuring 2 hidden layers of 16 neurons each. The lifts were in every case expanding windows as in equation (7.5). On top of this another recurrent neural network was placed, and the value of its final hidden state used as the output. This final network used 2 hidden layers of 16 neurons each.

E.3 Non-Markovian deep reinforcement learning

We used the implementation of the Mountain Car problem implemented by the OpenAI Gym [19], modified to return only the car's position. Each episode was run for 300 steps, and each model was given 2000 episodes in which to learn. The reward function was given by the car's position, in the range $(-1.2, 0.6)$, with a bonus $+1$ on reaching the goal. At each step the car could drive its engine left, right, or not use it at all. This problem was chosen for its ease of implementation.

The sizes of the models were chosen to ensure that they both had roughly the same number of scalar parameters. Within this specification, there was a small amount of hyperparameter searching. This was done in an *ad hoc* manner, for both models, varying the number of layers and the numbers of neurons in each layer, around the values that were eventually used. The eventual values chosen for the deep signature model were selected as the ones giving the best results for the deep signature model.

The eventual values for the RNN were selected to give roughly the same parameter count as the deep signature model, as no RNN model achieved any appreciable success.

The deep signature model was as described in Section 7.6, with the first network Φ^{θ_1} applying a learned linear transformation with output dimension 2. Furthermore it kept the original time-augmented stream, so that

$$y_i = \Phi^{\theta_1}(x_i) = (t_i, x_i, \phi_1^{\theta_1}(t_i, x_i), \phi_2^{\theta_1}(t_i, x_i)) \in \mathcal{S}(\mathbb{R}^4),$$

where $\phi_1^{\theta_1}$ and $\phi_2^{\theta_2}$ are learned linear functions. The signature was truncated at $N = 3$. The second network f^{θ_2} was comprised of a single hidden layer of 64 neurons, followed by an output layer of 3 neurons, corresponding to the three possible actions. The action with the greatest value was the action selected. This model had a total of 5769 scalar parameters.

The RNN model featured 3 recurrent layers each of size 32, followed by an output layer of 3 neurons, corresponding to the three possible actions. The action with the greatest value was the action selected. This model had a total of 5475 scalar parameters.

The reinforcement learning technique used was Deep Q Learning [111, 131], to effectively transform the task into a supervised learning problem, with each of the specified models performing function approximation on Q. Actions were chosen in an ε-greedy manner, with $\varepsilon = 0.2$. The discount factor was given by $\gamma = 0.99$.

The deep signature model achieved success, and would learn to consistently solve the problem at around 1500 episodes. The RNN failed to achieved success within 2000 episodes on any test run. 3 test runs were performed.